




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06K 9/36, 9/46	A1	(11) International Publication Number: WO 99/10838 (43) International Publication Date: 4 March 1999 (04.03.99)
(21) International Application Number: PCT/US98/17580 (22) International Filing Date: 25 August 1998 (25.08.98) (30) Priority Data: 08/918,791 26 August 1997 (26.08.97) US (71) Applicant: THE TRUSTEES OF COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK [US/US]; 116th Street and Broadway, New York, NY 10027 (US). (72) Inventors: BONNEAU, Robert, J.; Apartment 1B, 523 West 112th Street, New York, NY 10025 (US). MEADOWS, Henry, E.; Apartment 10K, 601 West 113th Street, New York, NY 10025 (US). (74) Agents: TANG, Henry et al.; Baker & Botts, LLP, 30 Rockefeller Plaza, New York, NY 10112-0228 (US).		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, HR, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the</i> <i>claims and to be republished in the event of the receipt of</i> <i>amendments.</i>

(54) Title: SYSTEM AND METHOD FOR PROCESSING DIGITAL IMAGE INFORMATION**(57) Abstract**

Color digital image data is encoded and compressed by a technique which combines the advantages of both wavelet and fractal encoding. The encoding technique divides the images (201) at numerous scales (270, 272, 274) produced by wavelet transformations and forms blocks comprising of groups of pixels at each scale. The average modulus values and average angles values of the data in each of the blocks are compared to the next higher scale. Each scale is then encoded for the blocks which have corresponding matching blocks in the adjacent scale.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

SYSTEM AND METHOD FOR PROCESSING
DIGITAL IMAGE INFORMATION

SPECIFICATION

FIELD OF THE INVENTION

5 The present invention relates in general to digital image processing and, in particular, to a system and method for applying a multiresolution transform to a black and white or color digital image for compression encoding and decoding of the image information, pattern recognition and video image processing.

BACKGROUND OF THE INVENTION

10 Digital image processing is an important area of advancement in the field of computer science with many current applications and an increasingly growing number of potential applications. The subject of digital image processing includes the storage, analysis and communication of images which are represented in the digital domain by a series of bits or bytes corresponding to each point in an image. A typical
15 example of a digital image is one that appears on a screen of a computer. The screen consists of a number of monochrome or colored picture elements ("pixels"), each of which have associated binary values which determine if the pixel should be illuminated (and in some cases how bright it should be illuminated). The simplest case is where each pixel has one bit of data associated with it on a black and white
20 screen. If the pixel is lit, then the value of the bit is set to one. If the pixel is not lit, then the binary value is set to zero. Each pixel could instead have a byte (8 bits) of data representing either the distinct color, particular shade of grey or some other information. A typical screen could have an array of 520 by 480 pixels to display an image. In order to store one complete screen containing an image where each pixel
25 has a corresponding byte of data to it, approximately two megabits of data would have to be used for this example (520 x 480). More pixels are used in higher resolution screens which are becoming more and more popular today.

In order to store a large number of single images in a database for storage and processing, a data compression technique is required to make managing the database efficient and feasible for operating in real time. In addition to on-site applications with digital images, digital images can be transmitted to an outside site either via a network, dedicated line or some other type of conduit of data. In order to increase the efficiency of data transmission and represent images which will fit in the bandwidth of the data conduit, the data must be also compressed. An imaging device for recording images such as a digital camera could be placed at a remote location, have the image data digitally processed and compressed at the remote location, transmit the compressed data to a central processing station or other final destination location, and decode the image information so that an operator at the final location can view the image. The decoded image could also be matched against a database of stored images for identification purposes. If the database contained many records of images to be matched, the images stored in the database would need to be compressed in order for the database to hold and process the required number of images for a particular application. Accelerated pattern matching may be required for potential applications such as identifying a criminal caught on a bank's videotape where batch processing for storage and transmission purposes of the matching operation could take up to several hours due to the vast size of the database.

While the compression of image information is necessary for pattern matching, some conventional compression techniques can lose important image information in the process of compressing the data. An important aspect of a pattern matching technique is to be able to preserve the essential features of an object, such as their edges. The physical differences in the objects of the images could be very slight and there may be many similar objects stored in a database to be distinguished and matched. An example is a database of people who work for a large company or live in a small town. The pattern matching technique could be used to identify persons at an entrance gate but would have to account for small difference in facial features in order to distinguish the people. The use of digital images of faces in a database is currently being used for storage. In New York State and other states, the pictures on driver's licenses are digital images which are stored and can be reproduced if a license

is lost. The next step is to match images of people captured on cameras at crime scenes to the driver's license database of physical images to identify the criminal. Digital images of fingerprints or other objects could also be used. Pattern recognition of images should not be limited to objects in the exact same position because objects are not always still, but the recognition technique should allow objects to be rotated and placed in any position when pattern matching.

Digital image processing also includes video processing. Video is basically a time series of single images (called frames). Each image frame when shown sequentially over time shows movement in the objects present in an image. Video image data can also be stored and replayed. One example of digital video images is the video clips that appear in popular software programs. These video clips can include clips from movies which have been digitally recorded or clips recorded by a camera and stored digitally in the computer. Video images can also be transmitted over long distances. One example is teleconferencing which shows the image of the speaker while talking at a remote location and shows the speaker's movement or expression.

Video images require a large amount of data to represent just a few seconds of video time. Each individual frame of the video must be stored and replayed to create a recognizable video image. Even if only a portion of the frames are stored, the sheer number of frames requires the image data be compressed. Video images can also be used in pattern matching schemes which could identify particular objects in the video images. This may allow an air traffic controller to identify planes if other communication systems fail.

From the above discussion, a digital image encoding scheme is desired which has a high compression ratio while still preserving the feature's important details such as its edges.

One compression scheme currently in use is called "fractal encoding". Fractal encoding takes advantage of the fact that many subparts of an image are repeated and therefore an image can be represented by a mapping of the portions of the image to only a fraction of the subparts of the image (called blocks). By mapping the image onto pieces of itself, a separate code book and word relating parts of an

image to other objects does not need to be stored. Fractal encoding subdivides an image to be encoded into blocks which taken as a whole make up the entire image. Some of the blocks may overlap and be different sizes. In conventional fractal encoding, the image is divided into two sets of blocks. The first set is the domain blocks which will be compared with second set of blocks called range blocks. The domain blocks can be rotated and have mirror images created in order to create more choices of domain blocks which can be compared against the range blocks. Each domain block is compared to each range block to determine the closest match. The mapping of the domain blocks to the range blocks is stored. Only information regarding matching blocks is used and the remaining blocks may be discarded thus compressing the data.

Fractal encoding does generate high compression ratios relative to other known compression schemes. A compression ratio is defined as the number of bits in the original uncompressed image compared to the number of bits in the compressed image. However, images which have been fractally encoded tend to produce blocky artifacts when reconstructed and decompressed. This is due to the data being organized in blocks. The fine edge information which is required by advanced pattern recognition systems is not satisfied by only using a block matching fractal encoding scheme.

Another technique for compressing digital image information is wavelet edge detection. Wavelet compression techniques exploit the fact that images have spatial and spectral redundancies which can be eliminated to reduce the size of the data structure used to represent the image. In simple terms, wavelets transform an image signal into a set of basis functions, much like the application of a Fourier transform which uses sines and cosines as a basis set. When the set of basis functions is applied, the original image is transformed into a set of coefficients. These coefficients can be further transformed when a derivative or gradient operator is applied to the basis set. The coefficients then take the form of edges in different frequency bands or scales which allows for an efficient means of image and video compression.

Wavelet transformations produce scales which decrease in resolution as the scales increase. The wavelet transform, when applied with a gradient operator, can remove texture from the image resulting in decreased reproduction quality. It would be beneficial to combine the compression qualities of fractal encoding with the shape preserving qualities of the wavelet encoding techniques.

Some techniques have been recently developed using aspects from both fractal and wavelet techniques. These techniques focus on taking fractal compression techniques which are traditionally applied in a spatial domain, and applying them in the wavelet domain instead. However, these techniques do not take full advantage of spatial similarities revealed by the gradient operator in the fractal portion of the technique, and thus lose image quality as the compression ratio for the technique increases.

SUMMARY OF THE INVENTION

In accordance with the present invention, there is provided a system and method for processing color digital image data by encoding the data to gain high compression while retaining important edge information, for decoding compressed image information which has been encoded, for matching objects within an image field to objects stored in a database and for encoding and decoding video digital images. The encoding method combines the benefits of the conventional wavelet and fractal encoding schemes in a unique way to take full advantage of both schemes. The encoding technique first spatially decomposes the image data from each color representation (red, blue and green) into initially two scales by a wavelet transformation. The wavelet transformation uses a quadratic spline basis set which enhances edges. At each scale, a low frequency and a high frequency image is generated during the wavelet transformation. The high frequency image thresholds out coefficients below a certain grey scale level. The high frequency point representations are then divided into blocks, where the higher frequency (lower scale) representations are called range blocks and the next higher scale blocks are called domain blocks. The average modulus and angle values of each range and domain block are then calculated and recorded. The gradient direction values are then sorted

independently for range and domain blocks and compared to find the closest values. If the closest match does not exceed a given threshold value, the block positions and modulus difference intensity and angle values are stored in a file to represent that portion of the image. If the closest match exceeds a threshold, another scale is used.

- 5 The unmatched domain blocks now become the range blocks in the next scale for the new domain blocks in the just created scale. When all the blocks have been matched at levels below the threshold, the process is complete. The low frequency image of the scale which had the last matching domain blocks is spatially decimated and stored.

- The encoded image can be decoded using a decoding technique in
10 accordance with the invention. First the low frequency image of the highest scale is spatially decoded. Then the low frequency image is fractally decoded using the high frequency image block relationship (modulus angle) and the low frequency intensity information. The low frequency image is then transformed to the next higher scale (lower scale number) with an inverse wavelet and the technique is repeated until the
15 image is reconstructed.

- The encoded data processed in accordance with this invention is very useful in pattern matching/recognition schemes. The stored data of the encoded data can be chain coded around the edges of objects in the image which helps identify distinct objects. The object can then be matched across the scales to determine the
20 hardest edges. Noise can be removed and the edge information can be compared to a database of encoded edge information for identification. The edge retention feature of the invention allows for precise matching and increased compression.

- Video encoding is a very efficient process when performed in accordance with the invention. Video is made of a series of frames, each of which is a
25 digital image. The first frame is encoded with the image encoding scheme of the invention. The optical flow is then calculated between the first frame and the next frame. The average optical flow of the range and domain blocks is then calculated. If the changes to the image are large enough for a particular range or domain block (found by comparing the average optical flow to a threshold), that block will be
30 recomputed to correspond to the new image portion. Only those portions of the image which have changed will be affected. The new domain and range blocks in

compressed form are then transmitted or stored to reflect the current state of the image being processed. If a large number of blocks are changed, the entire next frame will be encoded in order to minimize error in the image.

BRIEF DESCRIPTION OF THE DRAWING

5 Further objects, features and advantages of the invention will become apparent from the following detailed description taken in conjunction with the accompanying figures showing a preferred embodiment of the invention, in which:

Fig. 1 is a flow chart of the steps for encoding image data in accordance with the invention;

10 Fig. 2 is a graphical representation of the encoding process of Fig. 1 applied to an original image;

Fig. 3 is a graphical representation of the range blocks shown in Fig. 2;

Fig. 4 is a graphical representation of the domain blocks shown in Fig. 2;

15 Fig. 5 is a graphical representation of the matching step of Fig. 1 applied to the example in Fig. 2;

Fig. 6 is a graphical representation of the spatial decimation of the low frequency image of scale 2 associated with the example of Fig. 2;

20 Fig. 7 is an example of a file of compressed image data generated in accordance with the invention;

Fig. 8 is a flow chart of the steps used for decoding compressed image information in accordance with the invention without using specific high pass information;

25 Fig. 9 is a graphical representation of the decoding technique described in Figure 8;

Fig. 10 is a flow chart of the steps for performing pattern recognition using image data encoded in accordance with the invention;

Fig. 11 is a graphical representation of multiple objects which are processed with the pattern matching technique of Fig. 10;

Fig. 12 is a graphical example of range block centroid matching used in the pattern matching technique of Fig. 10;

Fig. 13 is a graphical representation of performing the pattern matching technique of Fig. 10 to an unencoded image;

5 Fig. 14 is a flow chart of the steps used to encode video in accordance with the invention;

Fig. 15 is a graphical representation of performing the video encoding technique of Fig. 14 to an unencoded image;

10 Fig. 16 is a graphical representation of multiple object tracking which is performed in accordance with the invention;

Fig. 17 is a flow chart of the steps used for decoding video in accordance with the invention;

Fig. 18 is a graphical representation of a system upon which the technique so the invention can be performed;

15 Figure 19 is a graphical representation of the data processing technique applied to color images; and

Figure 20 is an example of a file containing compressed color image data generated in accordance with the invention; and

20 Figure 21 is a table showing coefficient values for discrete wavelet functions used in accordance with the invention.

DESCRIPTION OF A PREFERRED EMBODIMENT

The present invention is directed to a system and method for encoding and compressing digital image information which achieves high compression, has selective and accurate feature preservation and is computationally efficient. Once the
25 image information is encoded and compressed in accordance with the invention, a further related technique is described which can closely reproduce the original image from the compressed data which could have been transmitted or stored. The encoding technique also allows for very efficient pattern matching of digitally represented objects within the image which is further described below. Finally, the encoding
30 technique can be adapted to video images for image compression and shape

recognition in the video images. The encoding scheme of the present invention combines elements of both traditional fractal encoding and wavelet encoding techniques in a unique way to take advantage of the strengths of both these techniques. The primary technique for encoding image information will be described first. Subsequently, the specific techniques all based on the encoding technique for decoding, shape recognition and video encoding will be described.

Figure 1 shows a flow chart of the steps involved to perform the image encoding technique in accordance with the invention. The encoding process compresses the data representing the image so that the information can be more easily transmitted or stored in a storage medium. The compression ratio currently achieved for the technique is 32:1 (every thirty two bytes of data can be represented by one byte of compressed data) with a level of noise of about (33.5 dB PSNR (peak signal to noise ratio)). The noise is the difference between the original image before encoding and the reproduced image. The data representing the compressed image (or identified objects in an image) allows for faster shape recognition because of its reduced size and allows for greater storage of compressed images which can be used for future pattern matching. The following technique for encoding image information is typically performed on a conventional representation of an image made up of pixels, or picture elements. An image field is the entire image being processed which may be made of numerous objects located on a background. Thus an image could be made of a 1000 by 1000 grid of pixels when 10% of the pixels near the center of the grid constitute an object. A desired image portion to be stored in an image field can be made of multiple objects such as three circles in a stop light. Therefore, an image of a stoplight will be made up of three circle objects and a rectangle object. An example of the encoding technique applied to a specific image will be shown in subsequent figures.

Step 101 in Fig. 1 spatially decomposes the image to be encoded into a first and second scale using a standard two dimensional wavelet transformation. Wavelet transformations will identify edge information by taking the derivatives of an (x,y) point of a smoothing function which will be applied to the image data of the image to be transformed and thereafter computing the modulus maximum (largest

amount and intensity of information at the point) which indicates the presence of edges in the image. The start of the encoding is described by the following equations 1 to 3:

$$\psi^1(x,y) = \frac{\partial}{\partial x} \theta(x,y) \quad (1)$$

$$\psi^2(x,y) = \frac{\partial}{\partial y} \theta(x,y) \quad (2)$$

$$\psi_s^1(x,y) = \frac{1}{s^2} \psi\left(\frac{x}{s}, \frac{y}{s}\right) \quad (3)$$

Equations 1 and 2 are the gradients of the smoothing function (θ) in either the x or the y directions.

The present invention uses a biorthogonal gaussian derivative spline basis set for the wavelet encoding. The biorthogonal gaussian derivative spline basis set allows for greater edge information to be retained by the wavelet due to its characteristics. The quadratic spline basis has not been previously used in a combined wavelet-fractal transformation. Most fractal-related wavelet techniques use a simple Haar basis set which is easy to implement in a fractal encoding scheme dealing with blocks as representations but does not retain great amounts of edge information. The Haar basis set consists of a square function in the frequency domain or block functions in the spatial domain. The use of a quadratic-spline basis set when combining wavelet and fractal techniques allows edge shape information in the wavelet domain to be better revealed for more accurate fractal block matching.

ψ_s^1 is the x derivative of the smoothing function at each scale s, where s contracts the function corresponding to a spatial decimation of space. ψ_s^2 is the y derivative of the smoothing function at each scale s. S is usually a power of 2. In the first pass of the technique, the gradient scale is two. This means the image will be spatially decimated by 2 in both the x and y direction. In each of any subsequent scales, the gradient scale will be the next power of two, i.e., 4 (2^2) for the second pass, 8 (2^3) for the third pass and so on. Next, the image f is convolved with the smoothing function where f is the function representing the image.

$$W_s^1 f(x,y) = f \otimes \psi_s^1(x,y) \quad (4)$$

$$W_s^2 f(x,y) = f \otimes \psi_s^2(x,y) \quad (5)$$

W_s^1 and W_s^2 are the wavelet transform functions in the x and y directions.

After computing the wavelet transform of the image to be encoded for a particular scale using the wavelet transformation function, the wavelet image will be divided into a number of sub-regions or blocks. The blocks will contain a certain number of pixels $N \times N$, usually a power of 2 corresponding to the wavelet scale. The modulus and gradient angle for each (x,y) pixel in each scale is first calculated. The calculations are computed as described in Equations 6 and 7:

$$M_s f(x,y) = \sqrt{|W_s^1 f(x,y)|^2 + |W_s^2 f(x,y)|^2} \quad (6)$$

$$A_s f(x,y) = \arg(W_s^1 f(x,y) + i W_s^2 f(x,y)) \quad (7)$$

The modulus is the amount of image power stored in the group of pixel for a given scale and gradient angle shows the derivative of an edge if present.

Step 103 then forms a point representation at the selected scales of the image based on those points whose modulus value exceeds a predefined value. For the first iteration of the technique, two scales will have point representations, designated scale numbers one (where $S=2$) and two (where $S=4$). The point representation will be used in the fractal portion of the technique. When a wavelet transformation is applied to an image, two representations of the image are created consisting of a high and low frequency representation. The first representation, termed high frequency image, contains all the pixels with modulus that exceeds a certain threshold. The low intensity modulus values correspond to the low frequency information and the high modulus values correspond to the edges, or more distinct features. The high frequency image information will have fewer data points because only those pixel exceeding a threshold will be retained. Thus any empty space in an image will be removed saving space in the data representation.

Step 105 then divides the high frequency image in the scale N into range and the high frequency scale from N+1 into domain blocks. The size of the domain and range blocks will affect the PSNR (peak signal to noise ratio) and compression ratio in the resultant reproduced image. The more domain blocks which
5 are generated, the lower the PSNR thus producing a cleaner image but the compression will be reduced. An effective quadtree segmentation is used to subdivide the high frequency image of the lower of the two scales into the range blocks since the wavelet basis set of the lowest scale includes all of the other scales. If the energy in the image data in a range block is greater than a predefined threshold level, then the
10 range block will be further subdivided so that the modulus in a particular range block will never exceed the threshold level.

The efficiency of the technique can be further improved if the range and domain blocks are only selected from the object image and the spatially decomposed object image in the different frequency ranges. Many objects have
15 naturally or designed repeating features and so restricting the block choices to representations of that object reduces computations during block matching and also reduces required storage of the data representing the object and thus minimizes necessary bandwidth for transmitting the encoded image. This restriction to "local" blocks also improves the efficiency of the decoding technique of Figure 7.

20 Some objects in the image to be compressed do not form a closed curve which can be used as the region from which the domain and range blocks are selected. In order to select the domain and range blocks for these types of objects, a straight line between the ends of the open curve is calculated and the area within the now "enclosed" object is used as the subset of blocks to perform the fractal encoding
25 portion of the technique. This allows self-encoding for each object in the image whether it is closed or not.

Step 107 computes the normalized modulus maxima and normalized angle for each domain and range block generated in step 105. At this point in the technique, the Hölder exponent, a calculation well known in the art, of the wavelet
30 transform can also be computed if desired for later pattern recognition. The Hölder

exponent will be explained fully in the pattern recognition section subsequently. The normalized modulus and angle values are computed by the following equations:

$$\hat{m}_{2j} = \frac{\sum_{k=1}^N \sum_{l=1}^N M_{2j}(k,l)}{\sum_{k=1}^N \sum_{l=1}^N A_{2j}(k,l)} \quad (8)$$

$$\hat{a}_{2j} = \frac{\sum_{k=1}^N \sum_{l=1}^N A_{2j}(k,l)}{Norm} \quad (9)$$

The calculated normalized values \hat{m}_{2j} and \hat{a}_{2j} are the average of the non-zero modulus or angle values, respectively, for a block at scale j . The "Norm" variable in equations 8 and 9 is the number of non-zero pixels in a given domain or range block. The normalized modulus and angle information is calculated for the range blocks of the scale N and for the domain blocks of scale $N+1$. The normalized average modulus and angle information are stored in a compressed image file which will be described in greater detail in Fig. 9.

Step 109 then matches the normalized angle and modulus values from each domain block in the image to be encoded at scale $n+1$ to each range block at scale n , where n is the current scale to be encoded. Thus the first iteration of the technique has the range blocks at scale one and the domain blocks at scale two. The second iteration would have the range blocks at scale two and the domain blocks at scale three. The average norm angle and the average modulus value for all the domain and range blocks are separately sorted by angle value and modulus value and then compared in a look up table. By sorting the normalized average values of modulus and angle, each block of the domain blocks does not have to be compared individually to each range block which is done in conventional fractal encoding. By comparing the modulus and angle values in pre-classified sorted order, a large savings in computing time can be accomplished which yields a sizable increase in efficiency of the encoding scheme.

Step 111 checks if the difference between the normalized modulus maxima and angle for a particular domain block to the closest range block is above a predefined threshold. The difference value is an excellent measure of how similar a domain block is to a range block. Because the modulus and angle values are sorted,

determining the closest range block is a relatively fast process. The differences are calculated by the following equations:

$$mdif = |\hat{m}_{2^{j-1}} - \hat{m}_{2^j}| \quad (10)$$

$$adif = |\hat{a}_{2^{j-1}} - \hat{a}_{2^j}| \quad (11)$$

If the minimum difference value between a particular domain block and the range blocks is above predefined threshold, then the range block does not match sufficiently to the domain blocks of the next higher scale, and another higher scale must be used for proper encoding for that particular range block. If at least one domain block has a high minimum difference value, then a further scale must be generated. If the difference value for a domain block is below the threshold, then the present scale is sufficient to compress and preserve the image to the desired level and the data for that range and domain block will be recorded in step 111.

Step 111 stores the domain block/matching range block locations in a table in a file along with the average difference value between the blocks and the normalized angle value of the range blocks for the range blocks which had a minimum difference value below the predefined threshold when compared. The intensity offset ("b") between the matching domain and range blocks corresponding in the low frequency portion is also saved for each domain-range block pair. An example of the table generated is shown and explained in conjunction with Fig. 7. The intensity offset is part of the wavelet transformation equation expressed as:

$$W_2^{j-1}f(x,y) = Q_L W_2^j f(x,y) + b \quad (12)$$

where W_2^{j-1} is the wavelet transform; $f(x,y)$ represents the image; Q_L is the domain to range block transformation; and b is the intensity offset. The values stored in the file will be a compressed representation of the original image which can be later decompressed after transmission or storage, can be used for efficient pattern matching, or can be used as part of video compression.

Step 113 checks if the minimum difference values between the domain and range blocks exceed the threshold for a given block. If it does, the technique

continues with step 115 to obtain range and domain blocks at higher scales which are similar. As the scale increases and resolution of the range decreases, there are fewer choices for matching blocks to be selected.

Step 115 spatially decomposes the image information at the next
5 higher scale. In this preferred embodiment, the scale will be increase by a factor of two. The low frequency image of the higher scale will then be spatially decimated by a factor of two. Thus if the first scale were decimated by two, the second scale would be decimated by four and the third scale would be decimated by eight. The scale determines the resolution and amount of spatial decimation of the image. The process
10 continues with step 103 where the range blocks are from the high frequency image which where previously the domain blocks and the domain blocks come from the newly generated scale. Step 117 fractally encodes the low frequency image of the scale whose domain blocks when encoded. The encoding is done with standard fractal encoding techniques. The domain blocks of the lowest frequency information
15 and domain blocks are matched together to allow for further compression of the wavelet representations. Alternatively, the low frequency image could be spatially subsampled in order to compress its image.

Figures 2 through 6 are graphical representations of applying the encoding technique described in Fig. 1 to an original unencoded image which is to be
20 encoded and compressed. Figure 2 shows an original image and the intermediate processing steps performed during encoding. In this example, three scales were generated during the multiresolution transform of the image. Labels 270, 272 and 274 show the scale numbers in the figure. Box 201 represents the original image before encoding. The image of box 201 in this example is a face with eyes, nose, mouth and
25 hair. The shaded portion of the image represents the low frequency information of the face which would be present in a normal image such as a photograph or drawing. Box 203 represents the low frequency scale "one" (first scale) image after the image has been wavelet encoded at the first scale. The wavelet encoding divides the frequency components in the image in half and generates low frequency scale one image 203
30 and high frequency scale one image 205. The high frequency scale one image contains the most edge information. The low frequency scale 1 image obtained from

the wavelet transformation has some of the low frequency information of the image preserved and some of the edge information (such as the hair). Some edge information is contained in all scales. Box 205 represents the high frequency image scale one after wavelet encoding at the first scale. The high frequency scale 1 image of the first scale will retain only those edges greater than a certain threshold. Thus noise or very soft edges with low modulus values will be eliminated.

The Zeroth scale captures more edges of the image than the other scales because it contains the most frequencies, and any edges will be retained in the encoded data in the first scale if possible. The second scale is a decomposition of the low frequency portion of the first scale resulting in a decreased resolution of the compressed image. Box 207 represents the low frequency scale 2 image after wavelet encoding at the second scale. The low frequency information of scale 1 is transformed using the wavelet function to produce the low and high frequency images in scale two. The frequency information of the original image 201 is still preserved but not to the extent in the first scale of the original image because decomposed representation of the original image is being transformed. Box 209 represents the high frequency scale 2 image produced after wavelet encoding at the second scale which still retains most but not all of the edge information. The edge information which is retained is not as complete as present in the first scale. The image representation in box 209 does not have the hair or mouth edges although it does show the other edges. The resolution of the edges in the image data of box 209 is less than the image data of box 205 of the first scale. The third scale is lower in frequency than the second and the resolution of the compressed image will decrease. Wavelet encoding the low frequency scale two image 207 produces low frequency scale three image 211 and high frequency scale three image 213. Box 211 represents the low frequency scale three image after wavelet encoding the image information from low frequency scale two box 207 at the third scale. At the low frequency, the low frequency information of the overall shape is retained but is less than the other two scales. Box 213 represents the high frequency image after wavelet encoding at the third scale. The edge information retained in the coded data is less than the other two scales and only the outline of the face is recorded.

The purpose of the multiple scales is to gain the benefits of compression and edge detection of the wavelet transformation information and to further compress the image using fractal techniques. The scales are used to help satisfy the condition in fractal encoding that each domain block have a similar range block based on the original image. However, by providing the different scales, the blocks can be matched across scales, where a domain block from a higher scale (and thus larger) is matched to a range block of a lower scales. In order to accomplish this, one additional scale must be produced for the highest scale of range blocks used. The domain blocks are always one scale higher than the range blocks in order to increase compression. Thus when the first scale in high frequency scale one box 205 is fractally encoded, the domain blocks must be derived from high frequency scale two box 209 of scale two. Once the average modulus and angles of the range and domain blocks have been calculated and sorted, the difference between the domain and range blocks are determined. If the difference for each domain block with the closest range block is below a predetermined threshold, then the domain and range blocks relative positions will be recorded in a file. Those encoded blocks are shown in box 224 in scale one. The encoded blocks for scale two are shown in box 226. If the difference between the closest range block to a given domain block is greater than the predetermined threshold, that particular domain block must be encoded at a higher scale. The domain blocks which were not recorded to a file are then processed and placed at random location at the next higher scale and a further higher scale is created to become the new domain blocks. Once all the domain blocks have been encoded and the difference value for each domain block is below a threshold, the highest scale containing range blocks, which are on a fixed grid, is itself fractally encoded to preserve low frequency information and to allow for the image to be reconstructed as explained later. In the example of Fig. 2, low frequency scale two box is fractally encoded to form encoded box 228. Alternatively, low frequency scale two box could be spatially subsampled to be compressed. The decoding algorithm which is explained in detail later will start with the low frequency image with the lowest resolution (highest scale) of encoded data containing the low frequency information

and add back in the edge information from the stored domain-range block relationships to ultimately form the original image 201.

Frequency graphs 250 depict one dimensional representations of the frequency components used in each scale of the multiresolution transformation in accordance with this invention. The image is initially transformed into the frequency domain using a basis function (e.g., a biorthogonal gaussian derivative spline basis) as part of the wavelet transformation technique. The original image is represented spatially as being in the entire frequency range which is represented as running from zero to f , where the frequency range encompasses the entire image. Scale one which is the highest resolution of the wavelet transform divides the scale by a factor of two into a high frequency scale one box 205 and a low frequency scale one box 203 as shown in graph 254. The low frequency range of scale one covers from 0 to $f/2$. The high frequency range of scale one corresponding to box 205 runs from $f/2$ to f . Scale two is decreased in resolution by a factor of two from scale one in this example. The low frequency area in graph 254 is now divided in half by the equivalent of low pass and high pass filter as part of the subsequent wavelet transformation to become a new low frequency image 207 and high frequency image 209 for scale two which is shown in graph 256. The low frequency range of scale two corresponds to box 207 and runs from zero to $f/4$. The high frequency range of scale two corresponding to box 209 covers from $f/4$ to $f/2$.

Scale three is then decreased in resolution by a factor of two from scale two in this example. The low frequency area in graph 256 is now divided in half by the equivalent of low and high pass filters to become new low frequency image 211 and high frequency image 213 for scale three shown in representation 258. The low frequency range of scale three corresponding to box 211 runs from zero to $f/8$. The high frequency range of scale three corresponding to box 213 covers from $f/8$ to $f/4$. Scale three would then be decreased in resolution by a factor of two to create a scale four in this example if another scale was required by the encoding technique. If a fourth scale was required, the low frequency component of graph 258 would be divided in half to form a new low and high frequency representation.

The relationship between the scales in group 250 show that it is possible to start with the highest number scale (i.e., lowest resolution) and iteratively reconstruct the high frequency scales until the image is reproduced. Thus when low frequency scale two image 207 is joined with high frequency scale two image 209 of scale two, the low frequency first scale image 203 will be produced. Low frequency image 203 can then be joined with high frequency image 205 of scale one to form the original image. The entire frequency spectrum of the original image can be recreated using this method. This dividing of the frequency ranges allows the technique to store only the data from the low frequency box of the highest scale recorded with the mapping between scales. In an alternate decoding technique described in Fig. 19, the high frequency image does not need to be separately stored because the information is part of the domain-range block mapping. The highest scale used for domain blocks is not needed to be stored because the domain information is stored compressed image file. The remaining low frequency boxes can then be sequentially recreated to generate the original image before encoding.

Arrow 220 shows the relationship between the range blocks in high frequency scale one image 205 to the domain blocks in high frequency first scale 209. The domain blocks for a given scale are mapped to the range blocks of the next lower scale. If there is no match for a given domain block in a particular scale, then the scale will be increased by one and a new matching pair will be sought. Arrow 222 shows the mapping between the range blocks in high frequency scale two image 209 and the range blocks in high frequency scale three image 213. Because the technique had sufficient matches for all the domain blocks in scale three to all the range, blocks in scale two and below, a further scale of range blocks was not required.

Figure 3 shows an example of how the quadtree segmentation in step 105 of Fig. 1 is used across each of scales to subdivide the high frequency images in each scale into range blocks. The quadtree segmentation shows how the range blocks will be allocated based upon the edges of the image provided. A pictorial description is shown in Fig. 3 of the transformed images at each of the scales which contain range blocks. Labels 308 and 310 show the scale numbers. Range blocks 304 corresponds to high frequency box 209 and range blocks 306 correspond to high frequency box

213. Range blocks 304 show the information in the high frequency box of the highest scale (lowest resolution) which contains range blocks. The image information is not as detailed as the lower scale (scale one) because of the multiple low pass filtering performed during wavelet transformation. The image is preferably overlaid with
5 blocks equal in area (although the size of the blocks could vary). Where an edge of the image is present in one of the blocks, that particular block will be present in the other scales of increasing resolution.

Range block 306 is shown for scale one and corresponds to box 205 of Figure 2. The range blocks are to be matched with domain blocks of the next higher
10 scales. The resolution of range blocks is increased by a factor of two. This means that overlaid grid will have four times as many range blocks then the higher scale and thus more information will be processed. The increased number of range blocks for the same image allows additional edges and features to be stored and represented than were found in the third scale range blocks 302. In particular, the eyes and nose of the
15 face of the original image are now represented by the range blocks 304 of scale two.

Figure 4 shows graphical representations of the domain blocks which have the high frequency images divided into a set of all possible domain blocks that occur at modulus edges. The image is referred to as a modulus image because only the blocks with edge information above a certain threshold are represented after
20 wavelet transform. If a block does not contain an edge because too little information is present, it will be ignored. A threshold level of the amount of image information is established to ignore domain blocks with a small amount of edge information or noise. The threshold level can be set at such a level which will increase the efficiency of the encoding technique while being balanced against the loss of edge and low
25 frequency information from removing too many blocks. Labels 405 and 407 show the scale numbers. Domain blocks 401 show a representation of only the domain blocks which contain edge information from the scale three wavelet transformation. All the remaining domain blocks have been ignored. Domain blocks 403 shows a representation of only domain blocks which contain edge information from the scale
30 two wavelet transformation. There is no scale one domain blocks because the domain blocks are always compared with the range blocks from one scale below.

Each individual range block of range blocks 304 and 306 and individual domain block of domain blocks 401 and 403 are then pre-classified by the average modulus and average gradient angle of the image information contained in each block. The equations for generating the classifications are detailed in the explanation of Figure 1. Labels 515, 517 and 519 show the scale numbers. Labels 521, 523 and 525 identify the block type or representation type. The average modulus value and average angle values for each domain and range block will then be sorted and stored in a compressed file.

Fig. 5 shows a graphical representation of matching the domain blocks from Fig. 4 to the range blocks of Figs. 3 in accordance with the steps of Fig. 1. Domain blocks 501 of scale three is matched to the smaller range blocks of the next lowest scale, in this case scale two. Domain blocks 505 of scale two is matched to the smaller range blocks 507 of scale one. The conventional way of matching a domain block to a range block is by comparing every domain block to a corresponding range block by least means square differencing, which is computationally intensive. However, in accordance with present invention, the average modulus and angle values of the image information in each range and domain blocks are stored and sorted in tables. The tables are then compared to see if there are matching range blocks to each domain blocks based on the average modulus and angle values. Once sorted, the entire list of range blocks does not need to be checked for each domain blocks, but only the pre-classified blocks with close to the same normalized average modulus and angle values. Thus a domain block with an low average modulus and angle will be checked against range blocks with low average modulus and angles. If the difference in values between the a particular domain block and the corresponding range blocks is greater than a certain threshold value, then there is not a sufficient match between the blocks for the given scales and another scale must be generated to further subdivide the image and check for mistakes. Generating three scales is a typical example of the required scales for encoding an image of picture quality.

There are a number of steps which can be performed to allow faster matching of domain blocks to range blocks. First, the number of domain blocks could be increased by decreasing the size of the domain blocks. The domain blocks could

be rotated or otherwise transformed to provide additionally matching options.

Moreover, the matching of domain to range blocks could be expanded to scales which are not sequentially related. For example, a range block is scale one could be matched to a domain block is scale three. The domain and range blocks can also be restricted to representations of the same object in order to improve the efficiency of the technique. These methods will increase the efficiency of the encoding process.

Figure 6 is a graphical representation of fractally encoding the low frequency representation at the highest scale with range blocks of the original image. In this example, the highest scale which contains range blocks is scale two and the low frequency image containing the image information is shown low frequency scale two image 601 (corresponding to box 207 of Fig. 2). Labels 609 and 611 show the types of blocks used for spatial decimation. The domain blocks 605 of the low frequency image 601 are then matched to range blocks 603 of the same image and are encoded using conventional fractal techniques. The mapping of the range blocks and the domain blocks which represent the low frequency image of the second scale are stored in the compression file. Alternatively, spatial subsampling can be used to encode the low frequency image.

Figure 7 shows a file organized in a preferred format for image data which was compressed and encoded using the technique described in Fig. 1. The object oriented storing format shown is useful is pattern recognition and video encoding. However, the file can be organized irrespective of objects if the application does not require object identification. Object oriented aspects of the invention will be described in subsequent sections. The data file 700 shown in Fig. 7 is organized into a number of columns. Column 701 in entitled "Range Block X" and contains the location of a particular range block relative to the X direction of a two dimension grid (X,Y). Column 703 in entitled "Range Block Y" and contains the location of a particular range block relative to the Y direction of a two dimensional grid. For example, if a grid has 100 points in a ten by ten array, the first block in the lower left hand corner would have coordinates (0,0), i.e., $x = 0$ and $y = 0$.

The range blocks shown in file 700 are chain coded such that the edges of a particular object are stored sequentially and if plotted would form the object.

Label 725 indicates a chain coded edge. For each object identified in the image, the range and domain block information of each scale which is used to encode the object is stored separately. In this example, range block and other information is shown grouped for the identified first object in the first scale in data 703; the information is grouped for the first object in the second scale with data 705; the information is grouped for the second object in the first scale with data 707; and the information is grouped for the second object in the second scale with data 709. Note that the data shown for each object would in actual use have many more entries. The number of scales stored depends upon the number of scales used in the encoding scheme.

Also stored in file 700 is the relative locations of the domain blocks for each object in column 721 entitled "Domain Block X" and column 731 entitled "Domain Block Y". Column 721 contains data of the domain blocks in the X direction of an (X,Y) two dimensional grid. Column 731 contains data of the domain blocks in the Y direction of an (X,Y) two dimensional grid. The identified domain blocks correspond to the range blocks identified on the same line of the file in columns 701 and 711. Column 741 is entitled "Average Norm Angle" and is the average normalized modulus angle calculated for the particular domain block. A domain block is made up of a multitude of pixels (Example could be 2, 8, 64, 256, etc.) and the average angle is calculated by the equations shown with respect to Fig. 1. The average block difference which is indicative of the average relative intensity of the pixels in the particular domain block (called "b") is stored in column 751 entitled "Average Block Difference". The three columns on the right of file 700 are used for video encoding and pattern recognition. These three columns will be explained in detail in when the pattern recognition and video encoding technique is described. Column 761 is entitled "Hölder Exponent"; column 771 is entitled "Flow V_x "; and column 781 is entitled "Flow V_y ". File 700 can also contain a header which includes information such as the highest scale factor of the compression (two in the examples of Figs 2-6), the number of objects in the image and the number of iterations use to encode each individual image.

Figure 8 shows a flow chart of the steps involved with the preferred embodiment for the image decoding portion of the technique in accordance with the

invention. The decoding process transforms the data representing the compressed image so that the coded and compressed information will be reconstructed to be a very close approximation to the original image. The decoding technique uses a cross scale approximation of wavelet coefficients to reconstruct the image. The more
5 iterations of the decoding steps performed for some portions of the decoding technique described below, the closer the reconstructed image will be to the original image. The following technique of Fig. 8 is for performing image decoding of a compressed image which has been encoded with the technique described in Figure 1 or a similar technique using a multiresolution transform. An example of the decoding
10 technique of Figure 8 applied to a particular encoded image will be shown in Figure 9.

The technique described in Figure 8 is different from the technique described in the parent application in that the high frequency information is not required to be separately reconstructed and added to the low frequency information in order to reproduce the original image. The low frequency information at the different
15 scales can be the only starting point for reconstruction using domain and range pairs based upon the high frequency images but performed on the low frequency image. The high frequency edge information is inherently recorded in the range-domain block mappings. By only applying the mappings to the low frequency information and using the intensity offset based on the low frequency scales, the technique can
20 obtain an image of acceptable resolution much faster by not requiring the separate generation and use of the high frequency information.

Step 801 upsamples the encoded low pass image information to begin the decoding operation. This step applies an inverse wavelet filter to the low frequency image information received (e.g., stored in a file in the form of Fig. 7). The
25 inverse wavelet filter expands the image by performing the opposite wavelet function used during the compression technique. If the wavelet function in the compression technique spatially decimated the image by a factor of four, then the inverse filter will expand the image by a factor of four.

Step 803 iterates using conventional fractal techniques to decode the
30 fractally encoded low frequency information which was stored in the encoded files containing the matching domain and range block locations for that particular image.

After a complete iteration is completed, the result will contain the image information for the particular scale. Following each iteration, the data is passed through a low pass filter in step 805. The fractal decoding for each scale will preferably be performed one or two times, which maximizes the speed of the technique versus the quality of the image.

Step 805 then processes the point scale representations of the image of the present scale to remove blocky artifacts created by the fractal decoding by thresholding. If the average modulus value of an image in a particular block is above a predefined threshold (and thus blocky and sharp), it will not become part of the image. The thresholding step removes the blocky features caused by the fractal encoding/decoding portion of the technique. This allows the advantages of wavelet transformations of edge detection to be combined with the fractal compression advantages. While step 805 is performed in the preferred embodiment, it does not have to be performed to complete the decoding technique. However, the resultant image will be clearer if step 805 is performed. Additional iterations of fractal decoding can be performed at a particular scale as needed but the image should be preferably passed through the low pass filter in step 805 after each iteration.

Step 807 applies an inverse wavelet filter (the same filter as applied in step 801) to the present scale to produce the next level of desired spatial representation (increasing the scale by one). The data is sent through an inverse filter to transform the wavelet encoded data for the low pass frequency information at the next lower scale. This process takes the information from the lower resolution scale (higher scale number) and creates an image whose frequency representation is in the next lower scale which has greater edge information and a higher resolution.

Step 811 checks if the decoded image is at the desired level of image resolution at the frequency current scale. This can be predefined by the program or can be assessed in real time as the image is being decoded. If the image is at an acceptable resolution, the decoding technique is done. For example, an operator may simply want to determine if an object appears in a normally blank image field. The operator would need a high resolution image for his/her purposes. If the reconstructed image is not at an acceptable image, then the technique continues with step 803. If

the resultant image at the last processed scale is not acceptable, the process in step 803 fractally decodes the image of the new image scale using the range-domain block relationship stored in the image file (e.g., Fig. 7). This information including the intensity offset, b , will provide further edge information for the image. Preferably one to two complete fractal iterations are performed to achieve the desired image at each scale. The image information is then passed through a threshold to remove blocky artifacts after each fractal iteration and an inverse wavelet filter is applied to transform the image into the next scale.

The only information initially required for decoding is the initial low pass representation in the lowest scale and the range-domain pair information and low frequency offset. The fractal and wavelet combination used for decoding then reconstructs the image to a close approximation of the original image. The reduction in information required to be stored from the prior art is an important advantage of the invention.

Figure 9 shows a graphical representation of the decoding technique described with Figure 8. The decompression technique is represented by the following steps:

- (1) $j = J$;
- (2) while ($j < J$);
- (3) $S_{2j+1} \times f(x,y) = Q_L \times 1/\lambda_j \times S_{2j} \times f(x,y)^* (\hat{H}_{j+1}, \hat{H}_{j+1}) + b$
- (4) $j = j - 1$;
- (5) endwhile.

These steps are illustrated graphically in Figure 9. Box 900 shows an inverse wavelet filler H_{N+1} which transforms the low frequency image information stored in an image file (see Fig. 7) to the next lower scale by expanding the image by a factor designated between the scales (e.g., 2, 4 or 8). Box 900 corresponds to step 801 of Fig. 8. Box 901 shows a representation of a digital image which has been encoded using the technique described by Figure 1 and has passed through the inverse wavelet filler in box 900. The function $f(x,y)$ has been processed by a multiresolution transform and the low frequency component of the compressed image at its lowest scale is represented as S_{2N} where N is the number of scales recorded during the

compression operation. Box 903 shows the low frequency component of the encoded image which is reconstructed as the lowest scale using conventional fractal techniques. The fractal function Q_{LN} 907 contains the range and block information which maps the high frequency portion of the current scale with that of the next lowest scale. b_N 909 is the intensity offset of the matched domain-range block information between the low frequency portions of the adjacent scales. The fractal reconstruction is performed on only the low frequency representation and the high frequency edge information is reconstructed into the image through the range-domain block mapping. The function representative of the image is normalized by the factor $1/\lambda_N$. After each fractal iteration, the information is passed through a low pass filter to remove block artifacts which fall above a threshold in the representation. After the low frequency image has been sufficiently reconstructed for the given scale with fractal techniques, the image information is passed through a recursive wavelet filter 911 to further restore the image. In the preferred embodiment, one iteration of fractal reconstruction at each scale is sufficient to reproduce the original image. However, additional applications of the fractal reconstruction and low-pass filter can be used to further increase the quality of the resolution.

The image is then processed at the next highest scale after being passed through the inverse wavelet filter. In this example the image information is processed at the second scale 913. The information is fractally decoded using range blocks and domain blocks using the function Q_{LZ} 917. The fractal decoding process uses the offset intensity value b 919 indicating the difference in intensities of the blocks in the low pass representations between the adjacent scales. Each iteration of fractal decoding at the scale 2 frequency is passed through a low frequency filter to provide thresholding and remove block artifacts. After a sufficient fractal decoding (typically one to two complete iterations) has been performed, the image information is passed through an inverse wavelet filter 921 (the same filter as inverse wavelet filter 900 and 912) and transferred to the next scale.

In this example, the image is then processed at the next lowest scale one 923. The image is again recursively fractally decoded using the range-domain block function 907 and the intensity offset 929. The intensity values b_1 929 from the

image file is used in the fractal decoding process to retain the low frequency information. The image information is then passed through a low frequency filter 931 to remove block artifacts and then through a reverse wavelet filter 932. The result of the decoding process is a very close approximation of the original image represented by $f(x,y)$ 933.

The decoding process described in Figures 19 and 20 allows for a fast and accurate decompression of the image information by starting with only the low frequency portion of the image at the highest scale and uses the range-domain relationship and offset information to construct the image retaining both the low frequency information and edge characteristics of the image with a minimum amount of information. The high frequency image information does not need to be stored separately which helps to increase the compression ratio and allows for faster processing.

Figure 10 shows a flow chart of the steps involved with the pattern recognition portion of the technique. The encoded and compressed data could be used only to transmit or store data for later recreation of the image, although the present encoding technique provides a powerful tool for pattern recognition. The range blocks are chain coded for pattern recognition which will identify the outside edges of separate objects. Thus if a bowl of fruit is the image to be analyzed, stored encoded images of a banana or other desired fruit can be compared to the images identified by chain coding the objects in the bowl of fruit. The pattern matching technique can be extended to identify any object of which an encoded pattern is already stored. Because the objects are encoded and compressed, the pattern recognition routines will be much quicker than if a conventional bit by bit match was attempted. Moreover, the edge information of the images to be identified is stored in accordance with the invention with better compression and easier matching capabilities.

Step 1001 of the pattern recognition technique encodes the image to be matched with the encoding technique described in Fig. 1. The result is a file listing of the relative locations of the identified domain and range blocks for each scale for an object as shown in Fig. 7. Step 1003 then chain codes the blocks by feature if not already done in the encoding steps. Chain coding is not required for simple storage or

transmission so it would not be performed in the encoding steps unless pattern recognition or video encoding was desired. Chain coding itself is well known in the art and stores the relationship of the edges of an object which helps identify the object. For each range block along an edge, both the relative position within the image and average angle and modulus are stored. The average angle represents the average gradient of the edges in the block, and the modulus shows the intensity of the image at that point. The chain coding continues until a complete image created from the edges is formed or the line of edges simply stops. If in a range block which should predictively contain edge information due to the modulus and angle values surrounding it but does not, that block can be corrected and replaced with the expected information. This may be determined by "neural network" or other decision making techniques known in the art. However, the end of an edge may signal the end of an object or a another object covering the first.

In order to determine if an object is "blocking" or overlapping another object in the image field, the Hölder equation, which is well known in the art, is used. The Hölder exponent records overall low frequency information of an object as well as the edge information. It is computed by taking the high pass image information of each scale generated during the compression technique and computing the slope of the wavelet coefficient's magnitude across the scales. The log (base 2) of the slope is then entered into the new image as an indication of low frequency information and edges. A negative Hölder value represents a hard edge and a positive value corresponds to naturally occurring soft lined phenomenon, such as vegetation. The Hölder exponent can be used in place of or in conjunction with the Lipschitz α for pattern recognition applications. The Hölder equation is the following:

$$|f(x) - P_n(X - X_0)| < C|X - X_0|^h \quad (13)$$

The Hölder exponent h of a distribution f at the point x_0 is defined as the greatest h so that f is Lipschitz h at X_0 , i.e., there exists a constant C and a polynomial $P_n(x)$ of order n so that for all x in the neighborhood of X_0 , equation (13) holds true.

Essentially the h criteria measures the intensity of the wavelet modulus as the image function progresses to successively higher scales (lower frequencies). If an object has a negative h exponent, the intensity of the wavelet modulus maxima (points of the image which exceed the threshold) stays relatively constant over a
5 number of scales and there is essentially a "hard edge" which stays the same from scale to scale whereas higher h indicate softer edges. Thus, edges with a calculated h can be used to characterize edges in objects in images which is an extremely useful property in removing noise from objects and for identification purposes. A negative h indicates occlusions where there are multiple overlapping objects in the image by
10 showing a drastic change in the image edge rather than a softer edge such as a transition from an object to a background.

Figure 11 depicts two objects in an image field which are to be matched to a stored image in the pattern recognition embodiment. The image data for each object will be stored and encoded by the multiresolution transform as described
15 in Fig. 1. The first image 1101 is shown as being partially obscured by the second image 1103. Using the Hölder equation to compute the h value for a given block along the edges of each image, the type of edge can be determined for each point in each object. The h value is based on how quickly the edge spreads out over the given scales, or on how consistent the modulus maxima value is for a given block as one
20 increases in wavelet scale. If the edge does not spread out, it is a "hard edge" and h is negative. This would indicate an edge created by occlusions because of the drastic sharp change in modulus value. If the edges do spread out over the given scales, then the edge is "soft" and h will be larger and closest to a value of one. A soft edge indicates that there is not an occlusion and the surrounding portions will not be
25 obscured by another object. If the edge is extremely soft (h almost equal to one), then the image information will have a natural characteristic such as vegetation.

Labels 1121, 1123 and 1125 indicate the h for an edge block of an object. In the example, block 1105 indicated has a relatively large h (close to a value of .9) for the point so it has a soft edge. Therefore, the stored edge should not be due
30 to another object blocking the one shown. Block 1107 has an h between 0 and .2 at the indicated point. Therefore, the edge is not as "soft" as block 1105, but is still high

enough to be considered the edge of the object and not a blocking object. Block 1109 has an h value between -1.0 and -0.5 and is therefore identified as a "hard edge". The edge is identified as an overlapping edge. An analysis of the modulus and angle values for the surrounding stored blocks in an object will identify which of the objects contain the edge in question without occlusion to complete the object. The remaining object which was partially obscured can then be matched for only that portion which is unobscured. File 1111 will contain the data for the objects chain coded and stored consistent with the file described in Fig. 7. File portion 1113 will contain the data for the first object and file portion 1115 will contain the data for the second object.

Step 1005 of Fig. 10 matches the image points across the scales using the Hölder exponent h from equation 13 described above in order to eliminate any noise for the object. Noise might consist of blurred edges or pieces of objects which should not be part of the image field. The calculated Hölder exponent h values are used to distinguish noise from the actual object information. If the h values are close to one (or a predetermined threshold), the edges will be very soft and the information will not indicate a hard edge of an object. Those blocks with the high Hölder exponent h values can be discarded as noise to create an image with higher resolution. Steps 1005 and 1003 can be performed at the same time after one h value is calculated for each block containing information.

Step 1007 preserves only those image parts which are consistent across the scales. This can be done by taking only those Hölder h values which have negative values or are within a specified range. The range may be from -1.0 to 0.5 . This will preserve only the clearly defined edges to be compared against stored images which are used to identify the objects in an image. The low frequency information portion of the image is not as important in pattern recognition as the distinct edges in the image. This step may be performed simultaneously with step 1005 where noise is eliminated.

Step 1009 then used a conventional candidate matching algorithm to identify objects in the image field when compared to stored objects. The candidate matching technique for matching objects calculates the centroid (center of mass) of the overall object and calculates the angle and magnitude from the centroid to each

block containing edge information for object. Fig. 12 shows the edge blocks of an image to be identified. Edge block 1203 is one of many blocks which contains edge information. The distance and angle between the centroid and each edge block is recorded in signature graph 1205. The signature graph will be the same for the object
5 1201 no matter how it is rotated or turned except for a phase shift in the calculation of the angle which can be adjusted for. The signature graphs of the image to be identified can be compared to signature graphs of stored objects to efficiently determine if a match is present. Alternative known matching techniques which can be used are neural network, eigenvalue or correlation matching.

10 Using the inventive encoding techniques, the objects have been encoded, compressed and transformed to the wavelet domain to preserve edge information using the lowest possible amount of data for storage. The compression feature allows many objects to be stored in a database archive which can be matched against to identify objects in a new image. Thus if the encoded compressed images of
15 every car model sold in the world over the last twenty five years was stored in a database, a system including a camera device to scan and store the image of cars could identify any vehicle which was scanned by the camera. Information of types of cars based on sticker prices, types of drivers and other information could be stored and processed with the images. Similar applications could include a database of stored
20 images of people who work at a manufacturing plant which requires high security measures. People whose facial images were not in the database could be quickly identified as outsiders which would alert company security.

After an image of an object has been matched to one in a database, descriptive information stored which is correlated to the matched stored image could
25 be displayed to help identify the object in the image field. A written description could be produced independently or as a text overlay on the image itself. If an object to be identified had been partially obscured, the matching technique would only be applied to the edge information associated with a portion of the particular objects stored in a database which correspond to the unobscured portion of the image to be matched.

30 Figure 13 is an example of applying the shape recognition technique described in Fig. 10 to an image 1301. The image is subdivided into a number of

scales through the process of encoding the image in accordance with technique of Fig. 1. There are three scales in this example. Labels 1330, 1332 and 1334 help identify the columns in the figure. Scale one corresponds to box 1307 which has the highest resolution. Scale two corresponds to image 1305 which is lower than the resolution of box 1307. Scale three has the lowest resolution and is shown in box 1303. When the edges are chain coded, the blocks which do not contain edges or have a small modulus value are eliminated because only edges over a specified threshold are chain coded as previously described. Thus image 1303 will be transformed into object 1309, image 1305 will be transformed into object 1311, and image 1307 is transformed into object 1313. The Hölder exponent h can be used to further define the edges of any objects and eliminate any noise. The resulting edges which have the α value within the desired range will be recorded in a compressed data file 1321. The compressed data file will have the same format as the file described in 7. For each block in a chain, the (X,Y) coordinate block position will be stored for the range and corresponding domain block. The average modulus difference between the blocks and the average gradient angle in the blocks will also be stored. Each object will have its own chain coded blocks as shown in object one portion 1323 of file 1321 and of object two portion 1325 of file 1321. The compressed data files for each object can be checked against a database containing the chain coded data for objects to be matched against. Both images will remain in their compressed form for the comparison. The pattern matching technique of centroid matching described with Fig. 12 can be used.

The encoding and pattern matching techniques can also be extended to video compression and video pattern detection. The motion of objects in a video stream can be modeled based on the optical flow of the edges in the detected imagery. The optical flow of an image function is defined by an intensity function represented by $I(x,y,t)$ has two components which are defined by the following equation:

$$\frac{\partial I}{\partial x} V_x + \frac{\partial I}{\partial y} V_y = -\frac{\partial I}{\partial t} \quad (14)$$

At a fixed time t , instead of solving the motion constraint in equation (14) for the image $I(x,y,t)$, the image can be smoothed with the smoothing function $\theta(x,y)$ dilated by a factor of $2j$. The smoothed image reduces the computational noise when estimating partial derivatives of finite differences and yields the following equation:

$$\frac{\partial I}{\partial x}(I_t \otimes \theta_j) V_x + \frac{\partial I}{\partial y}(I_t \otimes \theta_j) V_y = \frac{\partial I}{\partial t}(I_t \otimes \theta_j) \quad (15)$$

5 Equation (15) allows the technique to recover the normal component of the flow from the wavelet transform at the scale $2j$. Instead of computing this normal component at all points (x,y) of a video image, the normal component is computed at only the locations where the wavelet modulus is locally maximum (exceeding a threshold). This technique saves significantly in computational complexity over
10 traditional optical flow computation techniques.

Equation (14) is used in the present invention to perform video compression and detection by computing the average optical flow with each block which contains edge information exceeding a threshold. Using a block based technique in video compression yields two advantages. First, we can detect flow
15 changes within blocks and predictively estimate the positions of both range and domain blocks. Using this technique to update video, only blocks where significant changes occur require updating which allows drastically reduced required bandwidth needed to transmit the video images. Secondly, localized wavelet transformations can be performed within a block and thus localized flow for updating an image can be
20 accomplished by only changing the data content of some blocks for each new frame. The wavelet scheme in this technique also allows a pyramid reproduction scheme which allows the technique to transmit low resolution frames when bandwidth requirements dictate and then increase the resolution for higher bandwidth applications. Moreover, specific objects can be tracked through scenes with the
25 optical flow technique. The file format shown in Fig. 7 can accommodate the image compression, pattern recognition, and video compression. For video compression, the values of V_x and V_y would be included for each domain and range blocks. Once video

encoding starts, only those blocks that change can be transmitted, thus the wavelet coding and optical flow parts of the coding process can become background processes and thus not consume as much computational bandwidth on the video encoder processors.

5 Figure 14 shows a flow chart of the steps involved with the video encoding portion of the technique in accordance with the invention. Step 1401 codes the first frame in a sequence of frames with the image encoding technique described by Fig. 1. Video is made of a series of images which are projected in sequence to form the perceptions of movement. If the image is of a boy throwing a ball in a
10 playground, each image in the series of images will have the ball slowly changing positions as the ball moves while the background may not change at all. Thus only a small portion of the images in a video may change frame to frame.

 Step 1402 checks if any more frames are to be encoded. If more
frames need to be encoded, the process continues with step 1403. If not, the video
15 encoding processed is finished and the technique ends.

 Step 1403 reads the next frame and computes the optical flow between the frame encoded in step 1401 and the frame just read. The optical flow will indicate any movement of the edges of an object between the frames. This step checks the optical flow over the entire image.

20 Step 1405 computes the average optical flow within each range and domain block which has image information which has changed between the two frames. The average optical flow in a block will enable the technique to determine if any significant change has occurred on the image on a block basis.

 Step 1407 computes the new range and domain blocks which have an
25 average optical flow calculated in step 1405 above a predefined level. If the average flow is below the threshold, the information has not changed sufficiently to make a visual impact. Therefore, the image file does not need to be changed at this time. If the optical flow is above the threshold, the affected range and domain blocks will be replaced with new range and domain blocks which reflect the change in the image. If
30 an object is to be tracked, then all new range and domain blocks will be recorded in a separate file in order to store the complete motion of a particular object.

Step 1409 then transmits any range or domain blocks which have changed from the previous frame (and exceeded the predefined threshold in step 1407). Because all parts of an image frame does not change in every frame, only those particular range and domain blocks which did change will be transmitted to a video monitor or storage medium to show the localized motion. The block information will be transmitted beginning with the highest scale (lowest resolution) and increasing in resolution depending upon the available bandwidth of the transmission carrier.

Step 1411 checks if the number of range and domain block with a calculated optical flow which exceeded the threshold is above a second predefined level. If it is above the second level, then sufficient changes in the image field have occurred to warrant encoding the entire image field again instead of making partial charges. This will ensure that any noise in the smaller changer will not be compounded. If the threshold is exceeded, the technique goes back to step 1401. If the number of optical blocks which have changed is below the threshold, then the technique continues to process the next frame in smaller segments with step 1403. The video encoding ends when there are no more frames to process as checked in step 1402.

Fig. 15 is a simplified graphical depiction of the video encoding technique described in Fig. 14. Labels 1520, 1522 and 1524 show the frame number. First frame 1501 of a video shows a face with eyes, nose, hair and a frowning mouth. Second frame 1503 shows the same face except the mouth is no longer frowning. Third frame 1505 shows the same face except for a smiling mouth. These images can be compressed and transmitted to a different location using the video coding technique of Fig. 14. In practice, there would be many more intermediate frames showing the changing shape of the mouth.

Image 1507 shows a representation of the face in frame 1501 compressed using the technique of Figure 1 which is at a scale with the lowest resolution (high scale number). Labels 1517 and 1519 show the number of times the associated image is transmitted. In accordance with the technique described in Fig. 14, the entire low resolution image is transmitted only once for video frames

1501, 1503 and 1505 because the changes in the image are not substantial. Image 1509 shows a representation of the face image at a lower scale (medium resolution). Again, because the changes in the face were not substantial, the data describing image 1509 is transmitted only once. Image 1511 is a representation of the face image at the highest resolution (lowest scale). Only that portion of image 1511 which changes and has optical flow will be transmitted for each frame. The only portion of the image which will be transmitted is data for those range and domain blocks encoding the mouth of the face. Thus for frame 1503 and 1505, only the domain and range blocks correspond to the mouth of the highest resolution image will be transmitted.

10 Transmitting only the changing features of the image saves significant transmission costs and allows video consisting of many frames to be processed.

Figure 16 is a graphical depiction of multiple objects which are being visually tracked. A real life example of tracking objects is tracking two airplanes in the sky. The tracking portion of the video encoding technique corresponds to step 1407 in Fig. 14. First object 1601 is moving in the direction indicated by arrow 1602. Second object 1603 is moving in a direction corresponding to arrow 1604. As each object moves, the optical flow of the objects change. The optical flow of each object which appears in the image field is stored in a file 1605. The optical flow characteristics of object 1601 are stored for each range and domain block of the object in file portion 1607 and the optical flow characteristics of object 1603 are stored for each range and domain block at the object in file portion 1609. The format of the files is shown the right hand columns entitled "Flow V_x " 971 and "Flow V_y " 981 of Fig. 7.

Fig. 17 is a flow chart of the steps for decoding video images which have been encoded using the steps of Fig. 14. Step 1701 reads the encoded data for each frame which has been transmitted or previously stored. Step 1703 checks if the data is optical flow information from only a portion of the frame for the entire frame. This can be determined from either a predetermined bit value or the size of the data be processed. If the data is only from a portion of the image, then the process continues with step 1705. If the data is an encoded entire frame, the process continues with step 1707.

Step 1705 updates only changes domain and range blocks and decodes the images with this new information. Thus in the example of Fig 15, only the domain and range blocks encompassing the mouth of the face would be transmitted and changed in the currently displayed image. The resolution of the decoded frame would
5 depend on the system bandwidth which defines how many image scales can be transmitted and processed.

Step 1707 occurs when an entire frame is encoded using the technique described in Fig. 1. The technique for decoding an entire image described in Fig. 7 can be used in this instance. An entire frame is encoded when the amount of optical
10 flow information for a given frame exceeds a selected threshold (see step 1411 of Fig. 14). The video decoding continues for each encoded frame transmitted or being processed.

Figure 18 shows a system 1800 in which the present invention can be implemented. System 1800 contains three portions, video and image encoding
15 portion 1801, pattern recognition portion 1821 and video decoding portion 1831. Video and image portion 1801 preferably includes a camera 1803, a digitizer 1805, image memory 1807, and three specialized processes 1809, 1813 and 1817 each which have a respective local memory 1811, 1815 and 1819. A microprocessor (not shown) for running a series of instructions and distributing data among the three
20 processors is also included. The components are connected by conventional connectors and data buses.

Camera 1803 may be a video camera if video encoding is required or could be a still camera if only a single image encoding is required. However, a video camera could also be used to encode a single image representing either a single frame
25 or series of unchanging frames. Camera 1803 can be a color camera which encodes a color image by detecting three different color frequency spectrums: red, blue and green. The composite of the three images will produce a close approximation of the original image. The camera could be within the housing of the encoder 1810 or could be a remote camera connected by a connector or transmission equipment.

30 Camera 1803 is connected to digitizer 1805 which forms a digital representation of the image. The representation will be made up of a number of

pixels, the number depending upon the specific equipment used. The digitizer 1805 is connected to image memory 1807 which stores the image data for each frame captured by the camera 1803. The microprocessor (not shown) in the video and image portion 1801 is connected to all the components either through common
5 connectors or a databus in a conventional manner.

Video encoding portion 1801 shows three special processors 1809, 1813, and 1817. These processors are preferably dedicated to specific tasks to gain the advantages of parallel and pipeline processing. Processor 1809 is preferable dedicated to performing the wavelet transformations on image data. Processor 1813 is
10 preferably dedicated to computing the optical flow from one frame to the next. Processor 1817 is preferably dedicating to matching range and domain blocks in the fractal part of the encoding technique. The results of the encode image or video frame are sent via databus 1820. The databus, for example, could be a PCI, VME or similar high-bandwidth bus to suit different configurations. While three special processors
15 are described, the present invention can be implemented on any number of processors.

Pattern recognition portion 1821 includes a compressed image database and a separate microprocessor for performing the matching techniques. Pattern recognition portion 1821 could be located in the same casing as video encoding portion 1801. Bus 1820 is connected to pattern recognition portion 1821.
20 The compressed image database 1823 contains all the images and objects in their compressed form encoded in accordance with the present invention which are use to identify an object in a new image. The database 1823 can be large and can be stored on such storage mediums as magnetic tapes, CD-ROMs, or any other storage medium for large amounts of information. The processor 1825 will perform the matching
25 technique described in figure 10 including performing the Lipschitz and/or Hölder exponent computations on the image to be identified.

The results of the pattern matching will be sent via network 1830 to video decoding portion 1831. Video decoding portion 1831 could be located in the same casing as either the video encoding portion 1801, the pattern recognition portion
30 1801 or both. Video decoding portion 1831 includes a video monitor 1833 and a separate processor 1835 with other necessary components for performing the video

decoding and other functions. Monitor 1833 allows a user of the system to see the video (or image of a single frame) as it is decoded with any information from the pattern recognition portion about the image. Thus if a user is watching a busy highway and has stored the image data for all the models of the cars in the world, when a car image is recorded by camera 1803, the image will appear on monitor 1833 with a written description about the type of car after a pattern recognition operation has been performed in portion 1821. Monitor 1833 also allows a user of the system to view the entire decoded images which were compressed using the technique described in Figure 1. The image that appears on monitor 1833 can also be printed out using a conventional printer. The printout can identify the source of the original image if that information is included in the compression files described herein. Processor 1835 performs the decoding operation and other necessary processes.

Figure 19 shows a graphical representation of a multiresolution transformation of image information including color information. The color of an image can be processed, compressed and stored so that a color image can be decoded to reproduce the image. The image is stored as three separate components of the image with different color contributions: red, green and blue. Color images are conventionally separated into these three colors by the scanning or recording device. The color components are then later combined to achieve a full color image. The image information for a color image is typically separated into these three colors.

Referring to Figure 19, the color image 1901 is separated into three color segments (by filtering for example), red image portion 1903, green image portion 1905 and blue image portion 1907. The green color image is then transformed by operation 2109 to a luminance component with the equation:

$$Y (\text{luminance}) = 0.177R + 0.831G + 0.011B \quad (12)$$

Luminance represents the image in a similar manner to a black and white image. The other color portions will distort the edge and low frequency information. The luminance portion allows the technique to identify edges and sharpen the image for later reconstruction. Each image is then compressed separately by the technique described in Figure 1 and stored as compressed data in file 1910 of Fig. 20. Each

compressed color image portion is stored separately as compressed luminance image 1911, compressed red image 1913 and compressed blue image 1915.

When the compressed image information is to be restored, the decoding technique of Figure 8 is used to restore the color images to very close to their original state. Luminance portion 1911 is transformed back to green image portion 1905 by performing an inverse transformation function 1917. The inverse transformation function corresponding to transformation function 1909 is $G = 0.135B + 1.32L + 0.22R$. The result of the decoding is the three color image portions red image 1903, green image 1905 and blue image 1907. These three color images are then combined using conventional techniques to produce the original color image 1901.

Figure 20 shows the compressed data file 1910 representing the color image shown in Figure 19. The file contains a header 2014 and compressed image information for each object in the image area. Each object is further subdivided into the three color spectrums described in Figure 21. Object 1 indicated by box 2012 is saved as file portion 2013. The object 1 file portion 2013 has a header indicating the average low frequency information value of the image α_{m01} , the color of the file portion (red, blue or luminance), the (x, y) position of the centroid of the object, the rotation angle of the object (2nd momentum angle) and the number of range blocks stored in the particular color file. The color file portion 2016 stores information relating to the domain and range blocks for the particular object. The domain block positions (x, y) are indicated by Domain X and Domain Y, the range blocks position (x, y) is indicated by Range X and Range Y, the rotational angle of the range block is indicated by R_L (the range blocks can be rotated when determining the matching domain-range pair), the intensity offset ("b") discussed with respect to the decoding techniques is indicated for each domain and range pair; and finally the Hölder exponent indicates the average low frequency information of the domain/range block for the object is indicated by h (and can be used for fast object identification). The same type of information for the remaining scales for the luminance portion of the object are also stored in the file in portion 2018. File portions 2020 store information regarding the luminance portion of object 2 and file portion 2022 stores information

regarding the luminance portion of object 3. One of the objects could be the background such as the mountains or grass in an image. The low pass image 2030 is also stored upon which the fractal and wavelet conversions will take place. The corresponding image information portions for the red spectrum for each object is
5 stored in file portion 2024 (including a separate header) and the file portion for the blue spectrum is stored in file portion 2026 (also including a separate header).

The inverse wavelet filter's coefficients are stored in the actual decoding implementation which corresponds to the inverse to the wavelet filter used in the encoding implementation. Therefore, they do not have to be reflected in the
10 image file because they are used for all images which are encoded.

Fig. 21 shows a table 2101 of the preferred values of the coefficients for the discrete wavelet filters which can be used for both encoding and decoding of the image data. Column 2103 shows the coefficient number for the discrete representation of the wavelet functions. Column 2105 shows the values for the low
15 pass wavelet function used in the encoding portion of the technique. Column 2107 shows the value for the high pass wavelet filter used in the encoding portion of the technique. Column 2109 shows the values for the inverse low pass used for the decoding portion of the technique. Column 2111 shows the inverse high pass filter used in the decoding process in the parent application. The inverse high pass wavelet
20 filters are not required for the decoding portion of the technique of Figure 7. The foregoing coefficient values are used with the wavelet transform function of equations 4 and 5.

The foregoing merely illustrates the principles of the invention. It will thus be appreciated that those skilled in the art will be able to devise numerous
25 systems and methods which, although not explicitly shown or described herein, embody the principles of the invention and are thus within the spirit and scope of the invention as defined by its claims.

CLAIMS

1. A method for compressing color digital image data, wherein said color digital image data is separated into a plurality of different color data groups, comprising the steps of:
 - 5 transforming one of said groups of color data into a different selected color, wherein said different selected color is more indicative of edges in said image than said data group which is transformed;
 - spatially decomposing each of said groups of image data into high frequency and low frequency images at a plurality of scales of decreasing frequencies;
 - 10 fractally encoding said high frequency data for each said group responsive to said high frequency data at said scales generated for said group; and
 - storing separately encoded data for each said group data responsive to said fractal decoding and comprising information representing modulus and angle value relationships between portions of said image at said plurality of scales.
- 15 2. The method of claim 1, where said spatially decomposing step includes forming point representations at each scale for each said group.
3. The method of claim 2, said fractally encoding step includes the step of dividing said point representations of each said scale into blocks.
4. The method of claim 3, wherein said generated blocks are
20 restricted to be selected from an object representation in said image being encoded.
5. The method of claim 3, said fractally encoding step includes the step of computing the normalized modulus and angle values of each of said blocks.

6. The method of claim 5, wherein said fractally encoding step includes matching said average modulus and angle values from said blocks of each of said scales to said blocks of an adjacent scale.
7. The method of claim 1, wherein said plurality of different color
5 data comprises red, blue and green data.
8. The method of claim 1, wherein said different selected color is luminance.
9. The method of claim 8, wherein said transformed group of data comprises data representative of green.
10. The method of claim 1, wherein said stored group data is
10 processed separately for each color.
11. The method of claim 1, where said stored group data comprises luminance data, and said luminance data is transformed to green data prior to being decompressed.
12. The method of claim 1, where said storing data further includes
15 storing data indicating the color of the data.
13. The method of claim 12, further including storing data indicating the position of the object.
14. The method of claim 13, further including storing data
20 indicative of the low frequency information of said object.

15. The method of claim 14, where said data indicative of the low frequency information comprises a Hölder exponent responsive to said processing technique.

16. The method for processing compressed color digital image data
5 representing an original color image, wherein said image data has been spatially decomposed into a plurality of scales, wherein information representing modulus and angle value relationships between portions of said image at the plurality of scales and a representation of said image at the lowest scale has been stored, the method comprising the steps of:

- 10 a. spatially decoding said stored encoded representation;
 b. decoding said spatially decoded representation responsive to said stored relationship information;
 c. transforming said decoded representation in step b to a representation in a next scale.
15 d. repeating steps b and c until said representation closely approximates the original image.

17. The method of claim 16, wherein said decoding step b fractally decodes said image.

18. The method of claim 17, where said transforming step c uses a
20 wavelet transformation.

19. The method of claim 16, where said decoding step a uses a wavelet transformation.

20. The method of claim 16, further including the step of filtering said representation after each said fractally decoding step b is performed.

21. The method of claim 20, wherein said filtering removes block artifacts from said frequency information.

22. The method of claim 16, wherein said spatially decoding step upsamples said compressed image data.

5 23. The method of claim 16, wherein said relationship information includes block information and said decode step b uses an average angle value for said block information for decoding.

24. The method of claim 23, wherein said relationship information further includes data responsive to average intensity values of said blocks.

10 25. The method of claim 16, wherein said digital image data comprises color information, and said image is stored as a plurality of different color data groups, and method for processing compressed digital image data is performed for each said data group.

15 26. The method of claim 25, wherein said plurality of data groups are combined to restore said original image.

27. The method of claim 25, where one of said plurality of stored colors is luminance.

28. The method of claim 27, wherein said luminance data is transformed to a different color prior to being combined with said other of said
20 plurality of color data groups.

29. A system for processing compressed color digital image data to reconstruct said color digital image comprising:

a storage medium for storing information indicative of said compressed image which has been spatially decomposed into a plurality of scales, wherein said information comprises information representing modulus and angle value relationships between portions of said image at the plurality of scales;

5 at least one processor which repeatedly transforms said compressed image by fractally decoding and spatially reconstructing said stored information at a plurality of scales using said relationships, wherein said image information is repeatedly transformed until said compressed image approximates the original image; and

10 a display for displaying said reconstructed image.

30. The system of claim 29, wherein said processor utilizes average modulus and angle value of block representations of portions of said image during said decoding function.

31. The method of claim 30, wherein said processor further utilizes
15 data responsive to average intensity values of said block representations.

32. The system of claim 29, where said spatially reconstruction function is performed with wavelet transforms.

33. The system of claim 29, where said stored information in said storage medium includes data organized by objects in said image.

20 34. The system of claim 29, where said stored information in said storage medium includes data organized by color components.

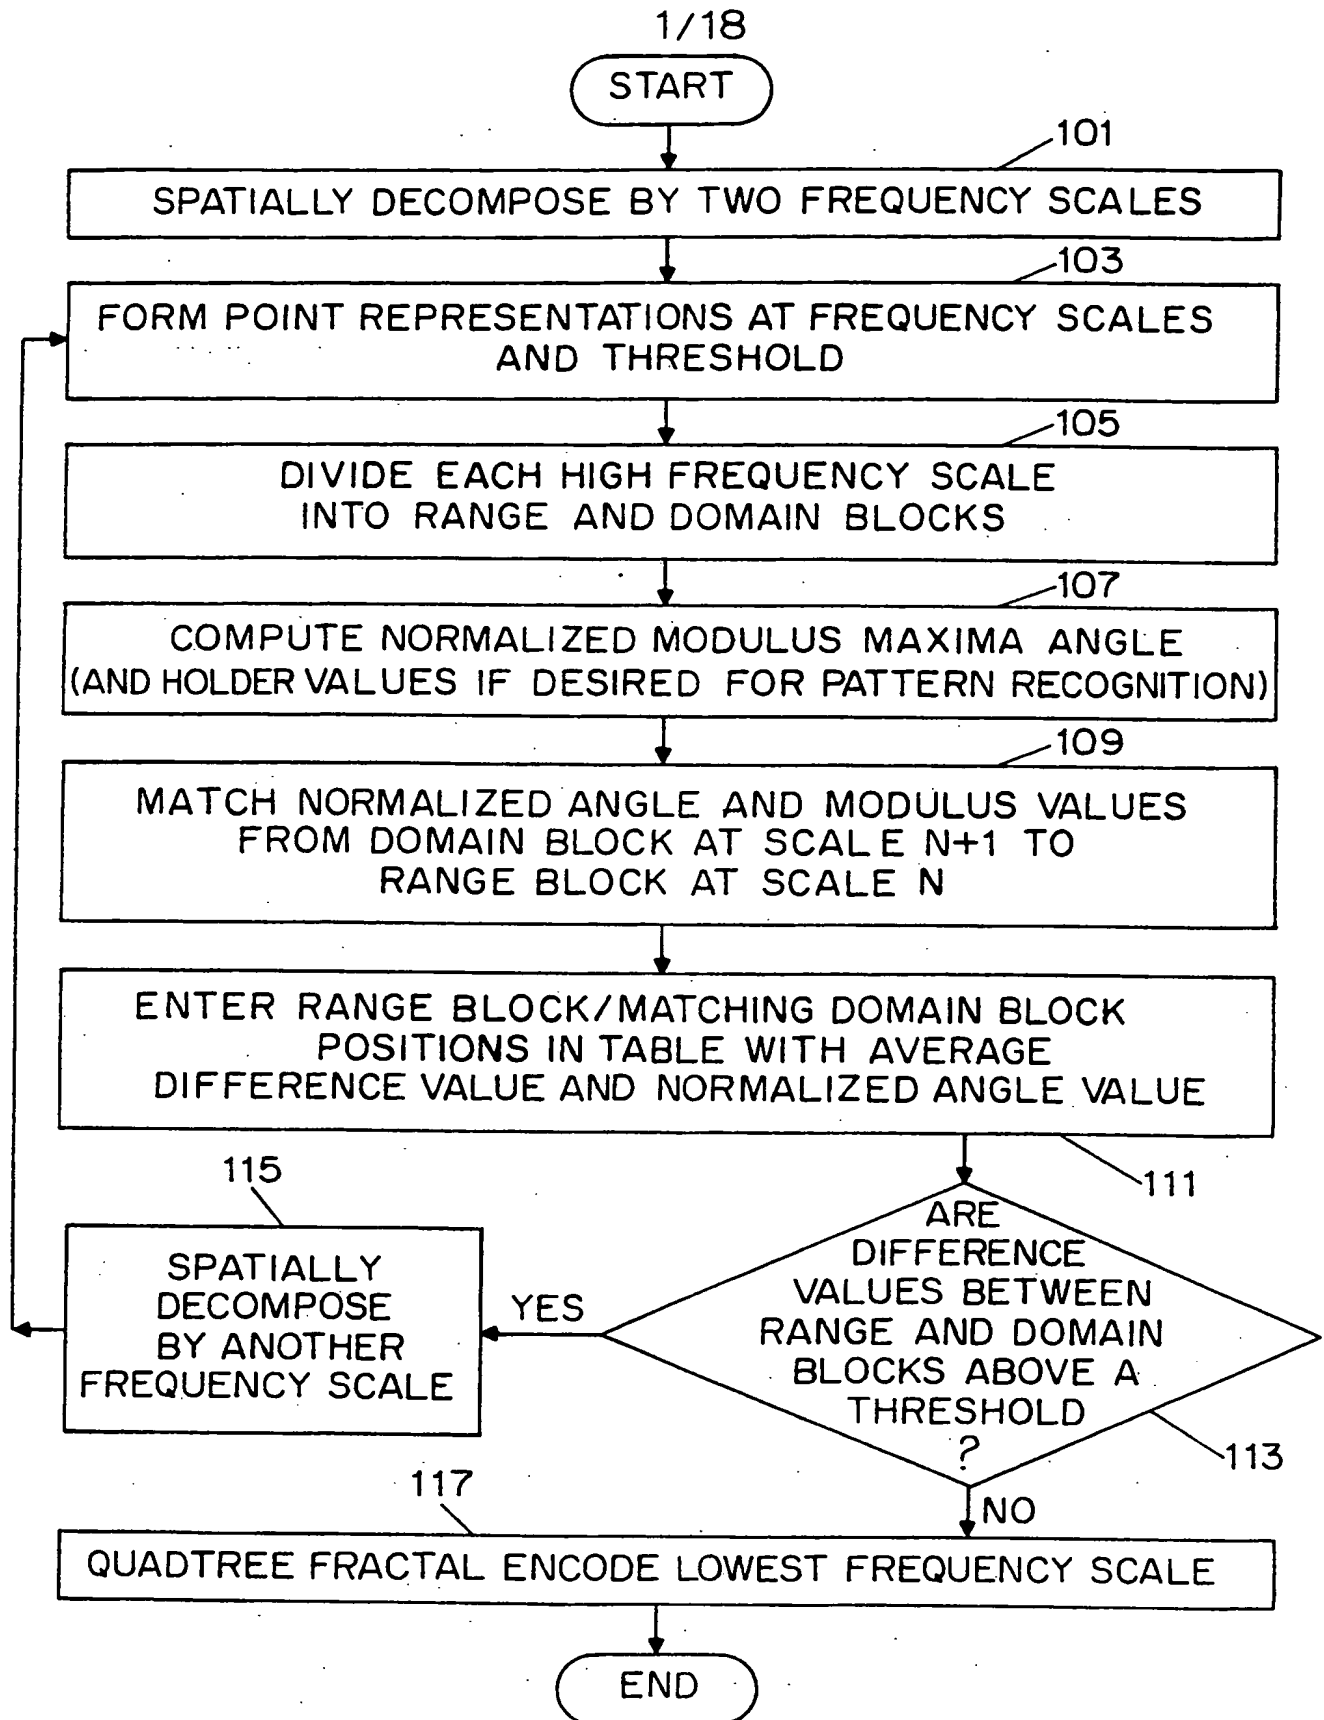


FIG. 1

2/18





FIG. 2

3/18




FIG. 3




FIG. 4




FIG. 5




FIG. 6

5/18

701	711	721	731	741	751	761	771	781
RANGE X	BLOCK Y	DOMAIN X	BLOCK Y	AVERAGE NORM ANG	AVERAGE BLOCK DIF	HOLDER	FLOW Vx	FLOW Vy
(b)								
OBJECT/FEATURE 1								
SCALE 1								
20	25	30	35	25	16	.15	10	10
15	20	31	40	12	12	.11	6	7
11	11	14	13	15	10	.12	4	4
SCALE 2								
20	15	30	35	•	16	.15	10	10
13	23	4	5	54	19	.01	14	15
11	11	14	12	15	10	.12	4	4
SCALE N								
OBJECT/FEATURE 2								
SCALE 1								
20	25	30	35	25	16	.15	10	10
15	20	33	40	17	12	.11	6	7
11	11	14	13	15	10	.12	4	4
SCALE 2								
20	25	30	35	•	16	.15	10	10
13	21	4	5	54	19	.01	14	15
11	11	14	13	15	10	.12	4	4
• • •								

703 705 707 709

725 CHAIN 725 CHAIN 725 CHAIN 725 CHAIN

FIG. 7

6/18

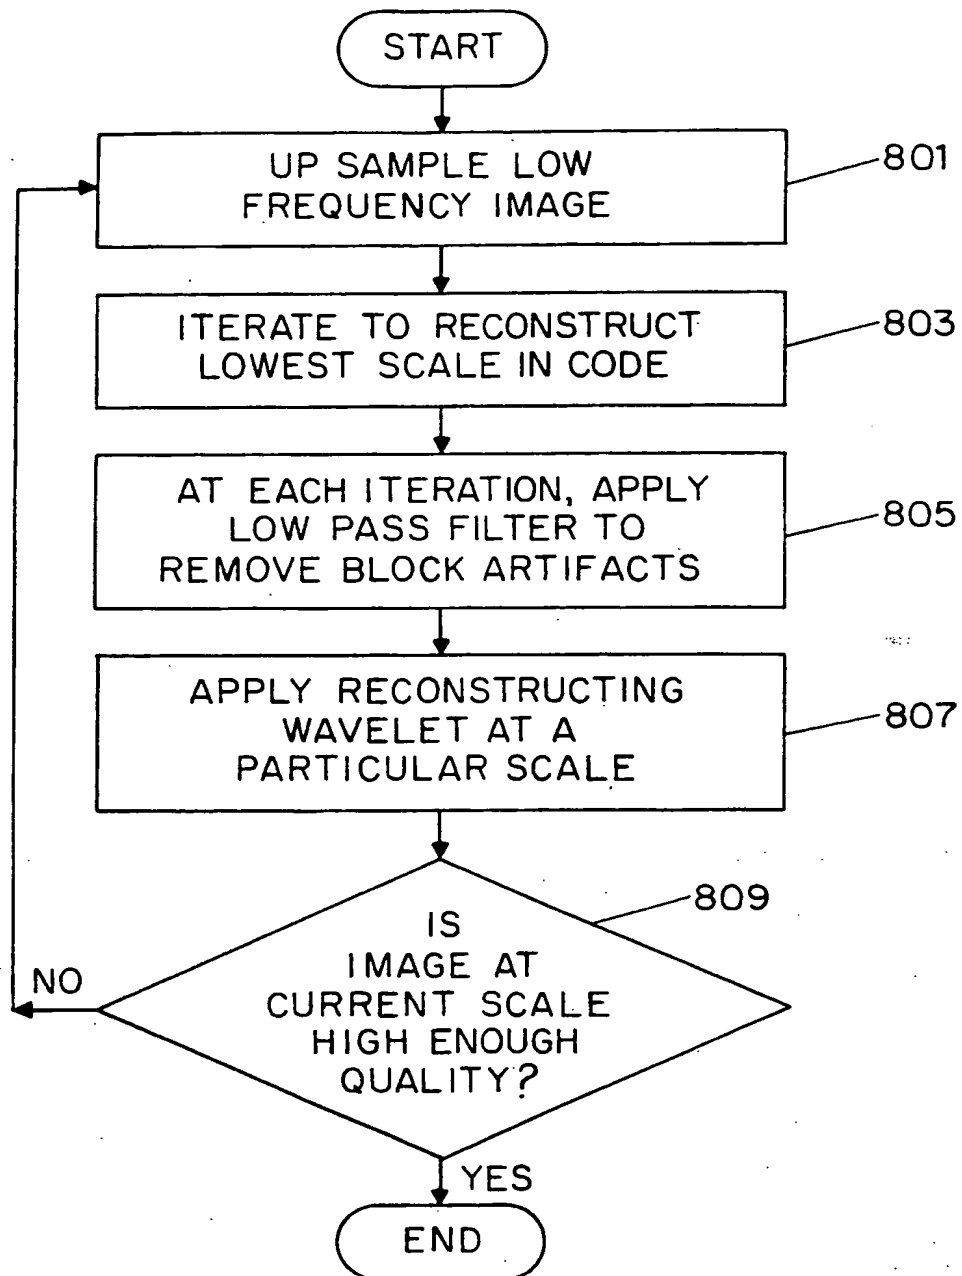


FIG. 8




FIG. 9

8/18

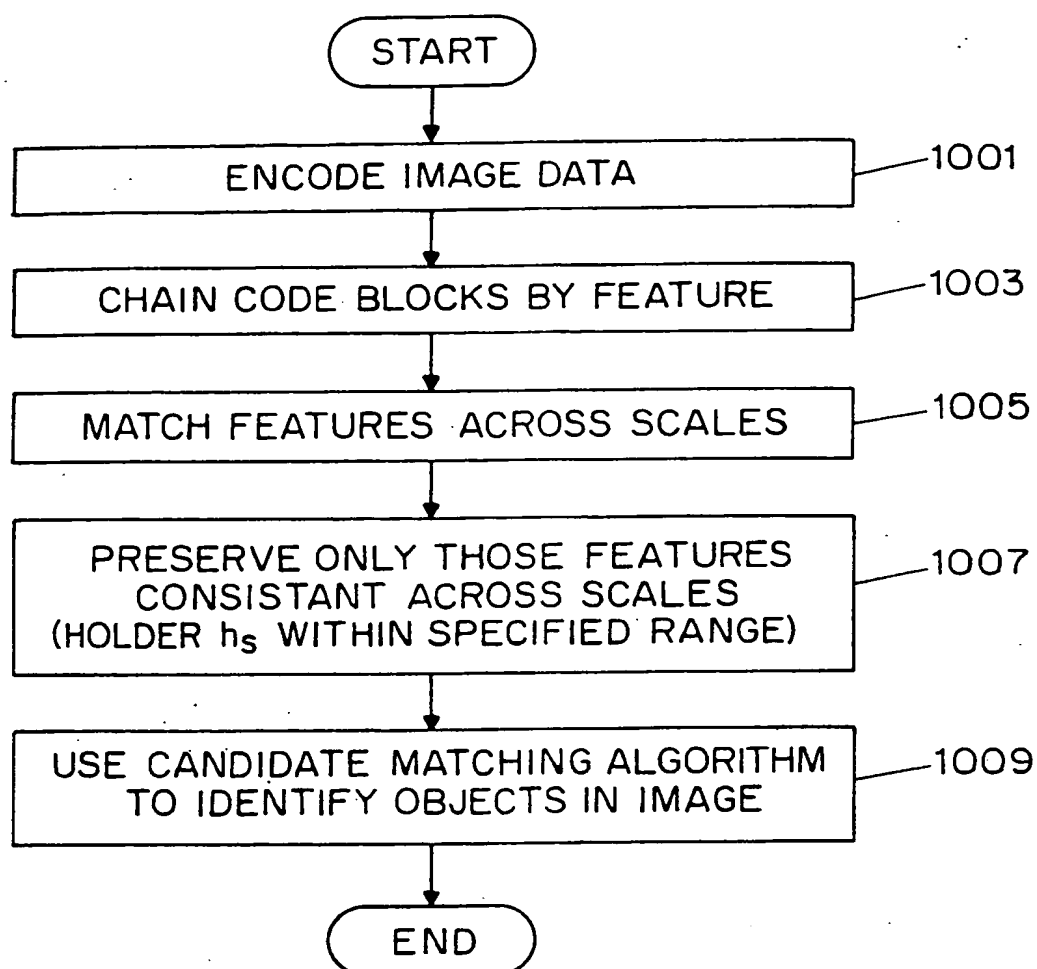


FIG. 10




FIG. 11




FIG. 12





FIG. 13

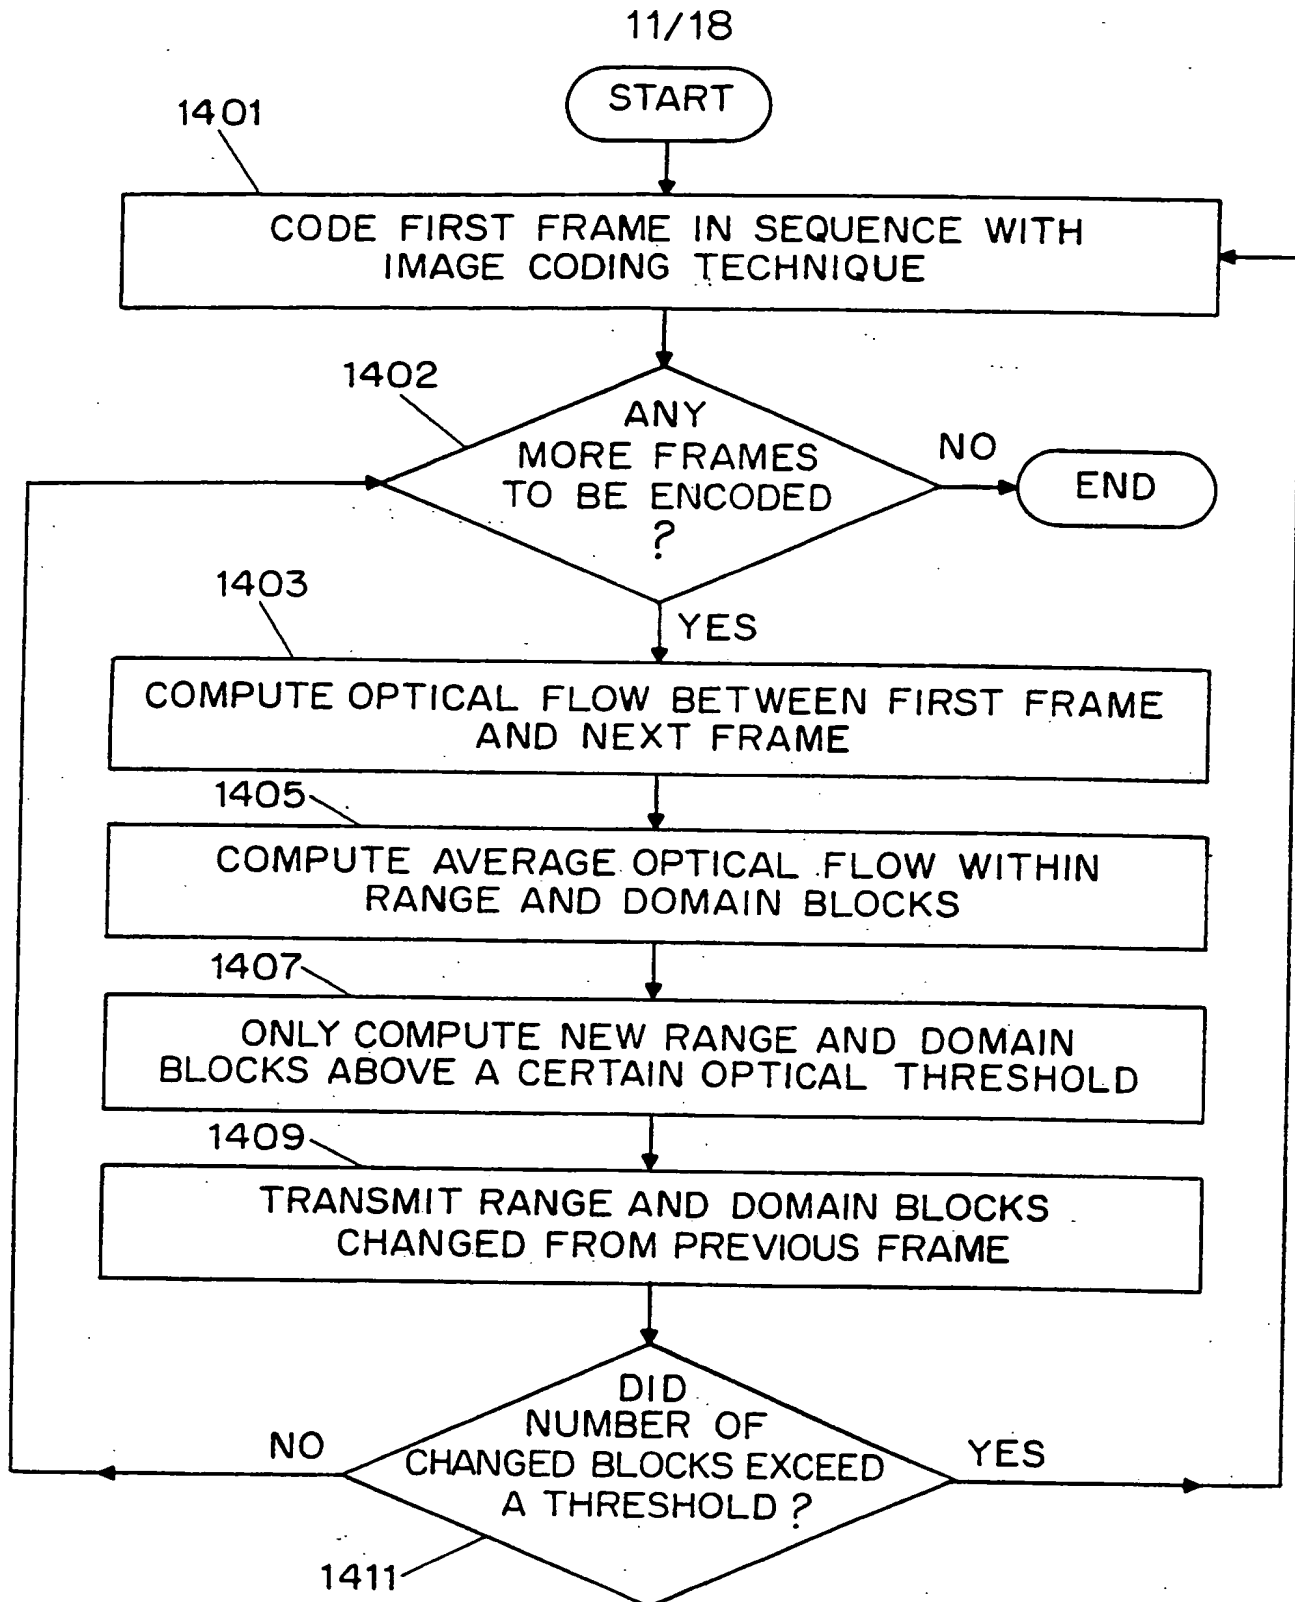


FIG. 14




FIG. 15

13/18




FIG. 16

14/18

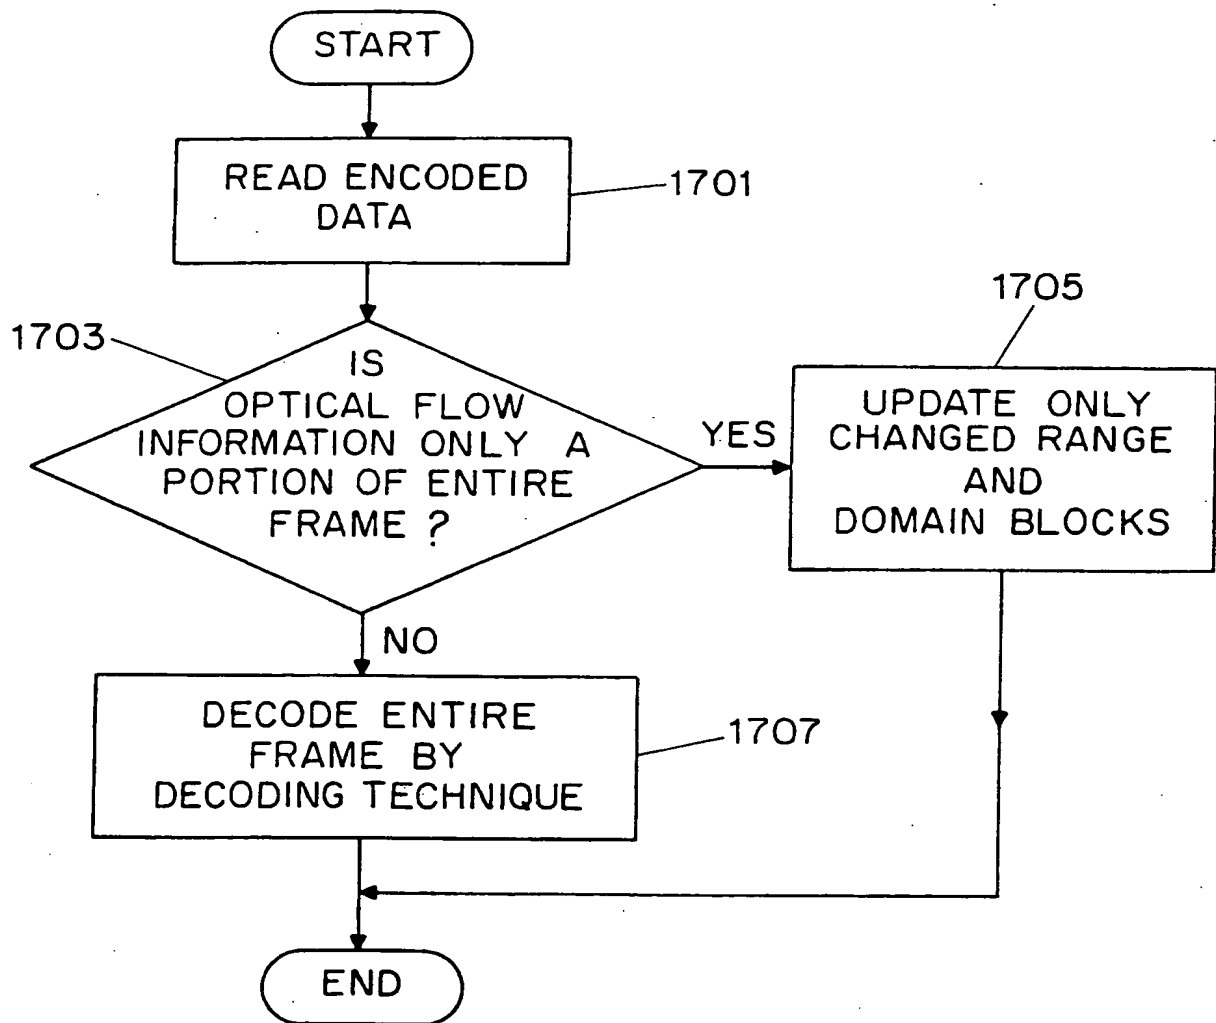


FIG. 17




FIG. 18

SUBSTITUTE SHEET (RULE 26)

16/18




FIG. 19

17/18

2012

2014

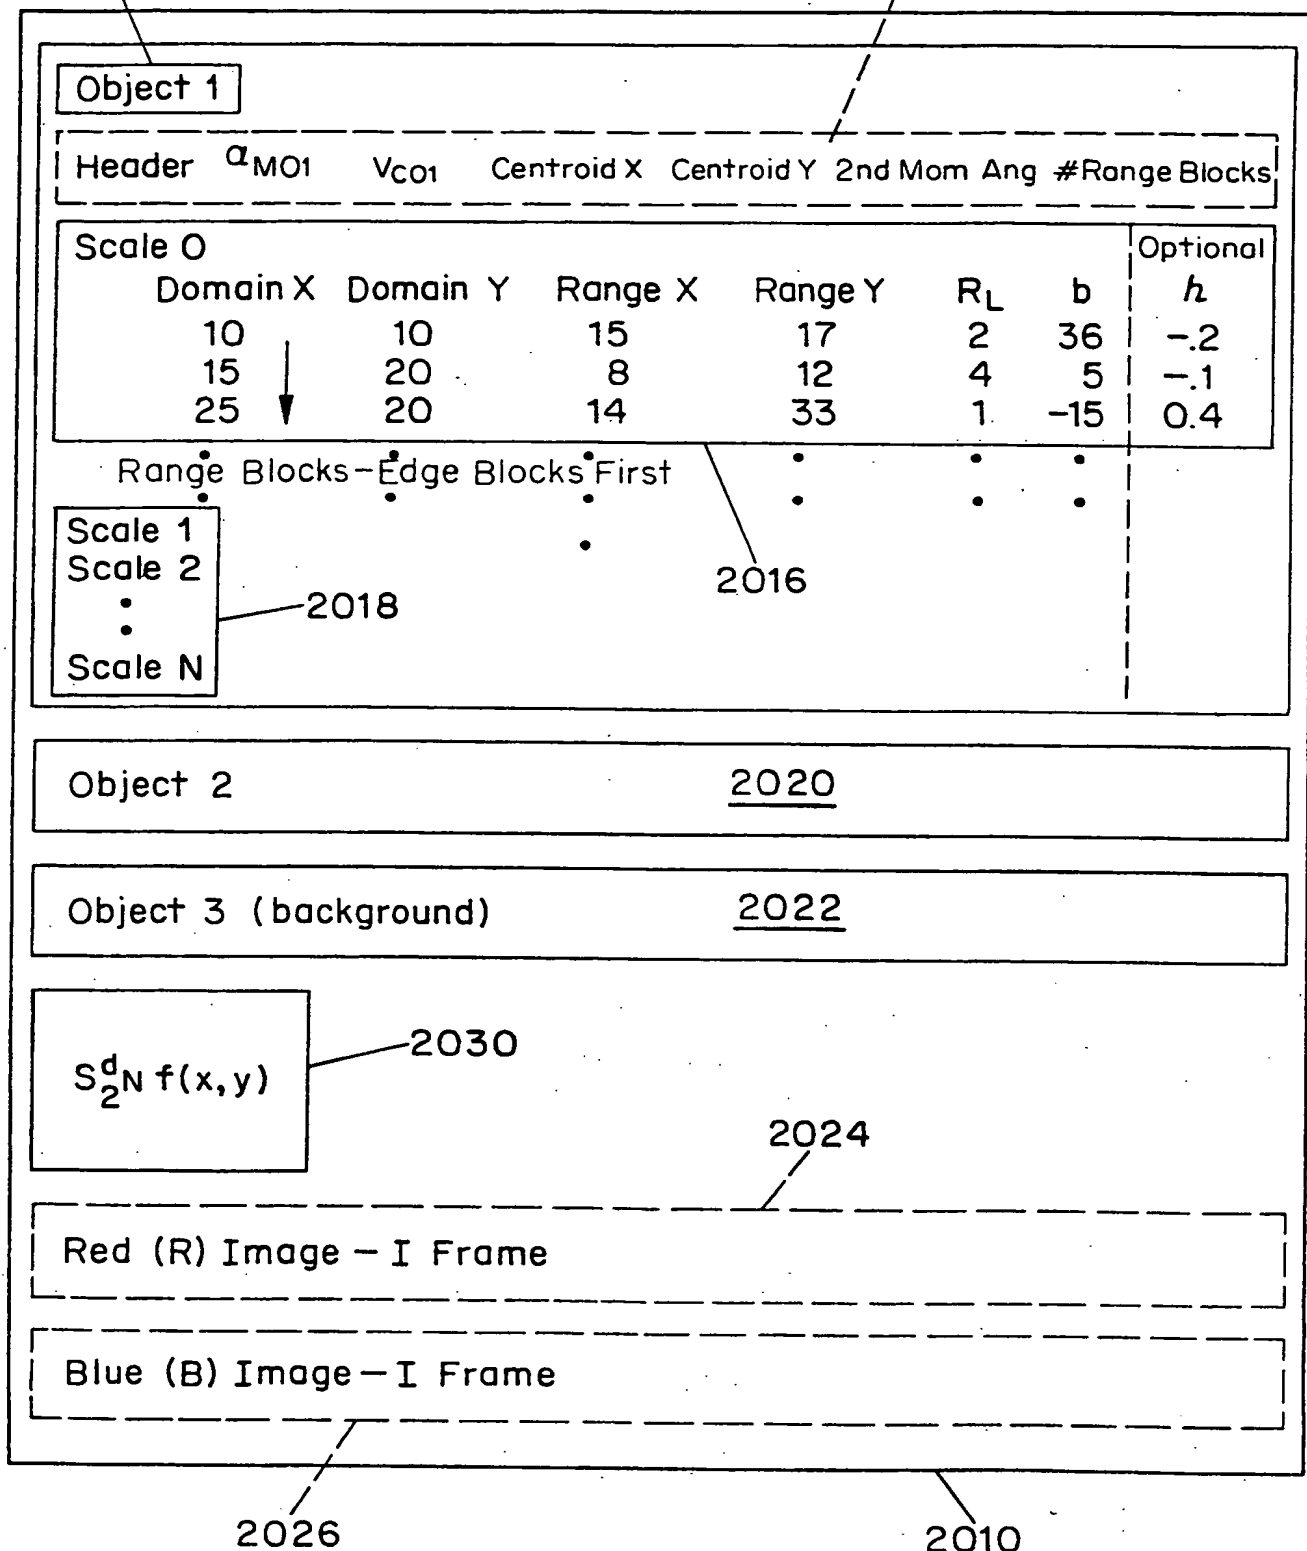


FIG. 20

18/18

2103	2105	2107	2109	2111
N	H	G	K	L
-3			0.0078125	0.0078125
-2			0.054685	0.046875
-1	0.125		0.171875	0.1171875
0	0.375	-2.0	-0.171875	0.65625
1	0.375	2.0	-0.054685	0.1171875
2	0.125		-0.0078125	0.046875
3				0.0078125

2101

FIG. 21